Ericsson Technology Review

#3, February 2026

Autonomous network operations: from reactive management to intent-driven optimization

Charting the future of innovation

Ericsson Technology Review — Charting the future of innovation | #3, February 2026

Autonomous network operations: from reactive management to intent-driven optimization

Authors:
Ciaran Johnston, Jörg Niemöller, Ann-Christine Eriksson, Wenfeng Hu, P.V.K. Ravikumar, Joseph Grogan

Traditional, reactive network operations are no longer sufficient to deliver reliable and cost-efficient telecom services. Autonomous network operations enable a paradigm shift toward networks that understand high-level business intents, self-orchestrate and continuously optimize to achieve desired outcomes. This evolution relies on artificial-intelligence-driven automation, robust observability and human oversight.

ISSN 0014-0171 284 23-3437 | Uen

Ericsson, SE-164 83 Stockholm, Sweden
Phone: +46 10 719 0000

The autonomous network (AN) vision aims to transform how telecommunications networks and services are deployed and managed to provide uninterrupted and optimized coverage, mobility and service continuity [1].

Traditional network management procedures are predominantly manual or follow deterministic rule-based automation, limiting the speed and flexibility with which a network can evolve, heal or optimize itself. In a shift to autonomous networking, intelligent machine-driven automation optimizes network capabilities dynamically to enable differentiated connectivity. The human role evolves toward observing and steering the automation through high level intents [2,3,4] and overriding machine decisions where necessary.

Intent and autonomous domains (ADs) are at the heart of the evolution toward AN operations. Intent expresses what the business or customer needs in measurable, policy-bound terms, while ADs define where that intent is realized safely – in bounded environments such as the radio access, transport and core networks – with clear guardrails, observability and control authority. This approach enables layered autonomy: independent domains that act quickly and reliably, coordinated by intent across domains to deliver end-to-end (E2E) service outcomes.

The architecture of an autonomous network

An AN is composed of one or more ADs, as illustrated on the left side of **Figure 1**. Each AD is built on a set of key components, including:

- intent management
- conflict management
- service-level objective (SLO) monitoring (as intent expectations)
- domain closed control loops (CCLs)
- domain intelligence (knowledge)
- artificial intelligence/machine learning (AI/ML)
- agent operations
- data management.

Together, these components enable faster fault resolution, proactive optimization and consistent Service Level Agreement (SLA) compliance, while reducing operational cost and human workload.

Intents are essential to achieving ANs, significantly simplifying network operations by enabling an operator to set clear expectations on the domain intent handler, with the intent handler taking responsibility for ensuring fulfillment of the provided intent. Intents are expected to propagate through different layers from business to resource, with more fine-grained intents at the lower layers. Potential conflict scenarios between competing intent goals are a key challenge that must be addressed.

Terms and abbreviations

AD – Autonomous Domain | **AI** – Artificial Intelligence | **AN** – Autonomous Network | **CCL** – Closed Control Loop | **CSP** – Communication Service Provider | **E2E** – End-to-End | **HITL** – Human-in-the-Loop | **KPI** – Key Performance Indicator | **ML** – Machine Learning | **NR** – New Radio | **RAN** – Radio Access Network | **SLA** – Service Level Agreement | **SLO** – Service-Level Objective | **UE** – User Equipment | **UX** – User Experience

SLOs are time-bound measurable targets defined as intent expectations that provide key insights about specific dimensions of service performance such as throughput, latency and jitter. SLOs can be computed for the purpose of evaluating service performance levels in the network. Communication service providers (CSPs) may offer connectivity services with contractual obligations in the form of SLAs. SLO observation is used to measure the fulfillment of the service performance according to the SLAs that the CSP delivers to its customers.

CCLs are the mechanisms that ensure accepted intents are fulfilled and expectations are met while the intent is active. CCLs resolve the observed delta between the current state and the wanted state of an intent expectation by aligning the former with the latter.

AI is an indispensable tool in achieving the AN vision. Various AI models are expected to interpret key performance indicators (KPIs) and predict future trends, preempting future service degradations and failures to improve customer experience and lower breach penalties. Agents and agentic systems [5] are designed to utilize AI and generative AI technologies to enable the development of automated on-the-fly troubleshooting and decision-making processes.

AI-driven network automation systems require access to substantial trustworthy and timely data to support use cases such as anomaly detection and failure prediction. Higher level autonomous systems rely on information (processed data) [6] to enable better outcomes in planning and optimization. Knowledge is further derived from domain- and/or application-specific information that adds to the accuracy of the outcomes and the explainability of decision-making by agents and language models.

Key challenges

The creation of an AN is heavily dependent on the capability of the underlying systems to support intelligence and automation. Managing multiple objectives is a significant challenge, as multiple stakeholders can place conflicting requirements on the same resource or service. Additionally, careful design of observability and explainability is required for operational users to be convinced to place their trust in an autonomous system, especially in the early stages, as it is learning the specifics of a particular network. Finally, ensuring that there is a smooth transition from the semi-automatic operating mode of today to the autonomous operating paradigm of the future requires a shift in the understanding of and a readiness to change the role of the network operations team.

Conflicting objectives

AN operations aim to fulfill multiple objectives that influence the same network services and resources. The objectives – usually expressed as intents – reflect the diverse interests of users, customers, regulators and the CSP in question. The challenge for an AD is to find the best possible outcome considering all the stakeholders.

An AD may consist of multiple software entities executing control loops that manage distinct concerns and pursue individual subsets of objectives. If not coordinated, such entities may issue contradictory operations toward network resources or subordinate ADs. One example is the Service Management and Orchestration framework with rApps [7] that are modular and control certain aspects of network operation. For example, network optimizations done by one rApp may have a negative impact on another rApp's attempt to optimize individual user service performance.

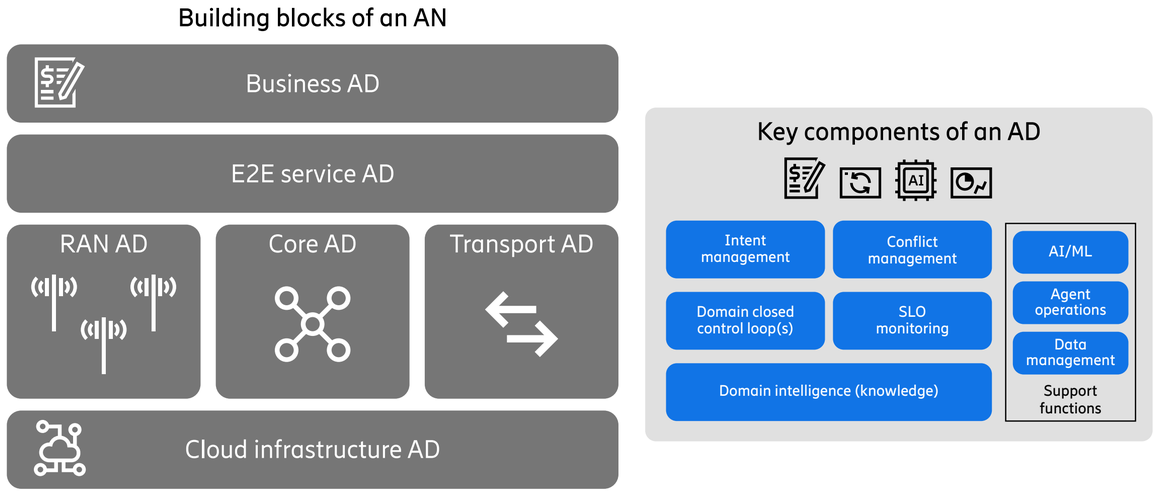

Figure 1: The building blocks of an AN (at left) and the key components of an AD (at right)

The intent handling for E2E service management will propagate requirements into multiple intent handlers across multiple network domains, taking care to perform feasibility and resource availability requests to avoid domain-internal conflicts. However, there are scenarios where not all potential conflicts can be mitigated at intent design time, requiring domain-internal handling of conflicts between multiple objectives.

Multiple objectives expressed as intents in AN operations can be seen as conflicts when it is not feasible to achieve all objectives. For example, fulfilling multiple connectivity service intents may lead to conflicts in situations where the available network resources are not sufficient to accommodate all the service requirements in terms of throughput and latency. Another example is when both an energy-saving intent and a connectivity service intent cannot be fulfilled at the same time.

In these situations, priorities and/or utility functions can be used to maximize the business value of the fulfillment of intents [3] by balancing the objectives within the control loops, utilizing necessary observability. This allows the CSP to express the relative importance of achieving different objectives to maximize the monetization of services and minimize network costs.

Inadequate management of control loops can lead to operational conflicts. Conflicts are direct if multiple control loops issue directly contradicting actions, for example, by

requesting different values of the same attribute in the network configuration. Direct conflicts are detected by analyzing and comparing the actions issued by multiple ADs and are relatively easy to detect.

Indirect conflicts occur when an action performed by one control loop causes a degradation in the fulfillment of objectives of another control loop. Such scenarios can result in cyclic behaviors. They are typically caused by collateral effects of actions and can be identified using impact maps. Impact maps are models that capture how actions affect KPIs not only for the target resource but also for all relevant associated resources. Impact maps can be constructed based on domain knowledge for qualitative estimates or ML models for quantitative predictions.

Figure 2 shows a simplified impact map, illustrating how three optimization actions interact in a small three-cell New Radio (NR) cluster, leading to indirect conflicts. Two actions are at the cell level: downtilt on cell 1 and increasing a parameter that adjusts the uplink power control behavior of user equipment (UE) on cell 2. The third action is at the inter-frequency cell-relation level: increasing the offset parameter that adjusts the threshold for switching cells to speed up UE handover from cell 1 to cell 3. Dark blue arrows indicate positive KPI impacts, such as better spectral efficiency on cell 1 or enhanced uplink throughput on cell 2. Orange arrows show negative side effects, such as reduced coverage from downtilt, increased load on cell 2, or extra traffic on cell 3 due to more aggressive mobility. Impact severity varies with radio frequency environment, user spatial distribution and temporal traffic patterns.

One way of protecting the network from direct conflicts is through authorization of AD scopes.

Managing conflicts between multiple control loops

There are a few methods for handling and resolving conflicts between multiple control loops. One way of protecting the network from direct conflicts is through authorization of AD scopes. In this scenario, actions in the network are uniquely assigned and enforced to a particular AD with no overlaps. The authorization is granted based on the function and role of the acting entity and its location in the network topology. This does not resolve conflicts, but it does protect the network from instability by enforcing clear rules regarding what actors can do. Furthermore, authorized entities may define the conditions under which other entities can act within the authorized entity's scope, thus including dynamic safeguards and policy enforcement in the AD scope authorization scheme.

Indirect conflicts can be handled by authorizing the scope of ADs such that they act on different timescales. For example, one AD could have a fast control loop that optimizes toward objectives by taking immediate action in response to KPI changes, while another AD may have a slower control loop that optimizes toward the same objectives with slower actions in response to KPI change trends over a set period of time.

Resolving indirect conflicts dynamically is primarily an optimization problem centered on finding the optimal network state considering all objectives and interests. Techniques based on CSP-defined rules and priorities can resolve most situations. A more self-adaptive and autonomous decision process can be achieved using utility-based decision-making [3]. Utility functions represent normalized information about the business value of outcomes and balance between objectives in prioritization decisions. In the scenario illustrated in Figure 2, an optimizer can systematically resolve conflicts by proposing the action set that provides the best overall trade-off across one or multiple intents by comparing a joint utility score for alternative action combinations.

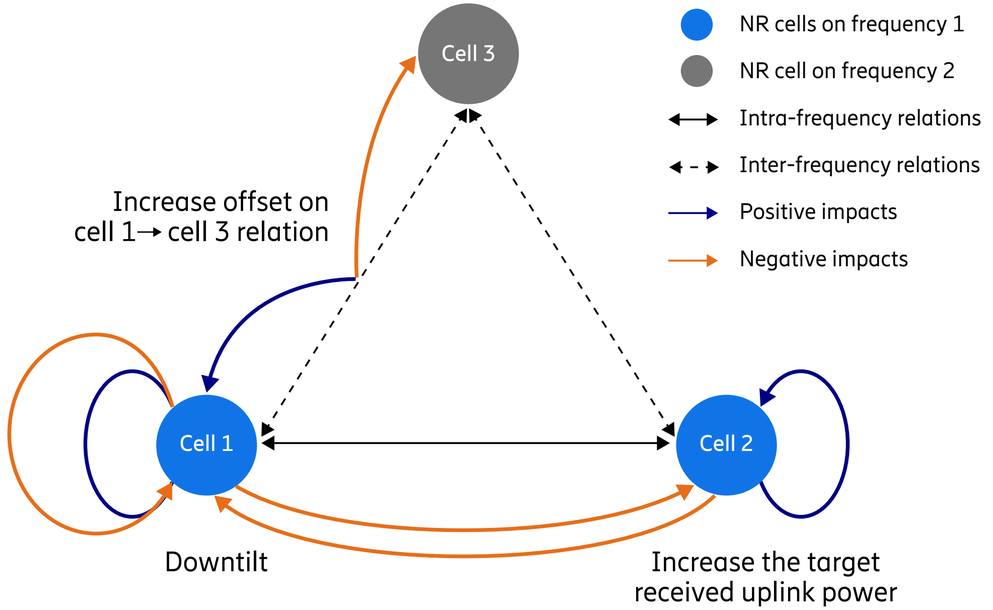

Figure 2: Simplified impact map illustrating three optimization actions with indirect conflicts

Collaborative conflict resolution

Further evolution of conflict detection and resolution can be achieved by establishing direct collaboration schemes between actors. If two actors understand the same type of actions, they can provide feedback and counterproposals to each other and negotiate on potential actions, with the final decision being made by the actor that has the authorization for the particular action. Another way to achieve optimized decisions across domains is for an actor that does not have authorization to act as needed to be able to send intents to the actor that has action authority. The decision-maker can consider the requirements from multiple ADs

and use all the features of the intent standards [4,8] such as feasibility checks, proposals and feedback, during intent negotiation.

Evolving the operator's user experience

Traditional network operations rely on teams of human operators supported by automation tools to maintain service quality, which means that while many processes are automated, human intervention and control remain essential for decision-making and execution. These teams are often reactive, responding to alarms and incidents after they occur. As networks evolve toward full autonomy, the operational paradigm shifts dramatically. The user experience (UX) for network professionals transforms from tool-centric, reactive workflows to intent-driven, proactive governance aligned with business objectives. This shift will entail significant changes to roles, interfaces, decision-making and organizational outcomes.

In semi-automatic environments, automation accelerates repetitive tasks such as configuration updates, fault isolation and software rollouts. However, these processes are typically triggered by events — alarms, performance degradation or outages — or are executed through carefully planned maintenance windows. Operators must interpret data from multiple systems, correlate symptoms and decide which scripts or workflows to execute. The UX in this mode is fragmented and operationally heavy, involving multiple dashboards and tools, performing manual correlation of telemetry, logs and alarms to identify root causes. It results in a reactive posture — most resolution actions occur after a service impact is detected. The cognitive load is high and success depends on human expertise embedded in scripts and tacit knowledge.

In an autonomous environment, decision-making shifts to intent-driven control.

Networks that are fully autonomous introduce a fundamental shift: operators no longer manage individual configurations or respond to alarms manually. Instead, they define business intents — desired outcomes expressed in declarative terms — and the system translates these dynamically into enforceable policies or configurations. Closed-loop automation ensures continuous compliance with these intents, using telemetry, analytics and AI-driven decision-making. At this point, the UX becomes goal-oriented rather than task-oriented. Instead of asking, "Which script should I run to fix this alarm?," the operator asks, "Is the network meeting the latency and security objectives for this service?"

In the transition from traditional to autonomous systems as shown in **Figure 3**, the interface exposed to control the network evolves to encompass:

- intent authoring and validation
- dashboards for continuous assurance based on actionable insights rather than raw telemetry
- explainability panels describing why the system acted as it did, which alternatives existed and what level of confidence it has in the solution.

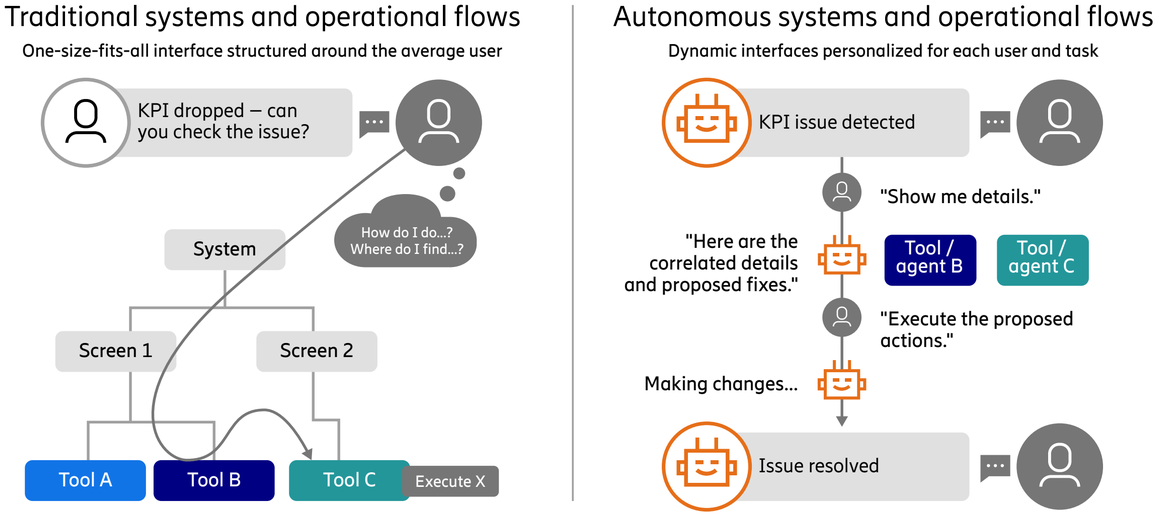

Figure 3: The transition from disjointed UX to collaborative interaction with autonomous systems

Key dimensions of the user experience evolution

The UX evolution from semi-automatic to autonomous operations has five key dimensions:

1. decision-making
2. the operational time horizon
3. context integration
4. the human role
5. risk management.

In semi-automatic operations, decision-making is highly manual. Operators trigger workflows and scripts based on their experience, often judging timing and scope under pressure. This approach places the cognitive load on individuals and creates variability in outcomes. In an autonomous environment, decision-making shifts to intent-driven control. Users specify desired outcomes and constraints, while the system plans, executes and continuously verifies compliance. Human interaction focuses on reviewing assurance proofs, confidence scores and risk indicators rather than parsing raw logs or telemetry.

The operational time horizon changes significantly in the transition toward autonomous systems. Semi-automatic models are reactive by nature, with troubleshooting dominating daily workflows and maintenance windows dictating when changes occur. This reactive posture limits agility and often results in service degradation before

corrective action is taken. Autonomous systems introduce predictive analytics and AI operations capabilities that anticipate issues before they impact customers. Dashboards evolve to highlight forecasts and risk trends and to recommended mitigations, enabling proactive intervention and reducing unplanned outages.

With respect to context integration, in traditional environments, context is fragmented across siloed tools for configuration management, telemetry and ticketing, requiring manual reconciliation to form a complete picture. This fragmentation slows decision-making and increases the likelihood of errors. Autonomous systems can consolidate these data sources, as well as external context such as news, weather and traffic data into a unified knowledge graph. This holistic context enables consistent decisions across domains such as the radio access network (RAN), transport, core and cloud, and provides operators with a single source of information for assurance and optimization.

The human role in the system changes substantially as well: the role of network professionals evolves from executing scripts and debugging failures to higher-level governance. In an autonomous system, operators become intent designers, policy governors and assurance stewards. Their responsibilities include curating guardrails, validating autonomous behavior and ensuring compliance with business objectives. This shift elevates human expertise from tactical troubleshooting to strategic oversight, aligning operational activities with organizational goals and regulatory requirements.

Governance must ensure that humans remain integral to decision-making.

Finally, risk management transitions from manual checks and slow rollouts to automated safeguards. Semi-automatic operations rely on human judgment to ensure safety, often through conservative change windows and rollback plans. Autonomous systems embed risk controls directly into workflows through guardrails, canary deployments and automated rollback criteria. The UX exposes these safeguards transparently, allowing operators to understand the boundaries within which autonomy operates and enabling them to intervene when necessary.

Trust and governance in an autonomous network

While autonomy does not eliminate human oversight, it does redefine it. Trust in autonomous operations is built on transparency, control and continuous assurance. As networks progress toward higher levels of autonomy, governance must ensure that humans remain integral to decision-making, especially in critical or high-risk scenarios. Three key capabilities are necessary to achieve this:

1. a human-in-the-loop (HITL) approach
2. dynamic observability
3. user-controlled autonomy.

An HITL approach is foundational to good governance: autonomous systems must provide mechanisms for human intervention at any stage of the decision cycle. HITL ensures that operators can:

- approve or override actions before execution in sensitive domains
- inject corrective intents when business priorities shift
- pause or roll back autonomy during anomaly detection or compliance audits.

HITL is essential for maintaining confidence and regulatory adherence, especially in environments where service-level guarantees and security policies are non-negotiable. Dynamic observability also plays a key role. Traditional observability relies on having many static dashboards and predefined metrics across many tools, each giving a unique perspective on the network. In an autonomous world, observability becomes dynamic and contextual – operators no longer hunt for data across siloed tools, but rather express prompts (natural language or intent-based queries) to retrieve correlated insights. Dynamic observability delivers explainability artifacts (why an action was taken, what alternatives existed) alongside predictive indicators (risk scores, drift alerts). The shift toward dynamic observability reduces the cognitive load of human operators and accelerates root-cause analysis without requiring deep tool expertise.

User-controlled autonomy is also essential – autonomy should never feel like a closed system. Operators must retain control over when autonomy is activated, with support, for example, for tiered autonomy levels (Observe → Recommend → Execute with approval → Full auto). The ability for the operator to define limits on autonomy to prevent cascading failures, as well as to steer and reprioritize the autonomous goals, is required for safe operations. This control framework ensures that autonomy is introduced in a progressive, trust-calibrated manner.

Conclusion

Autonomous network (AN) operations transform telecom networks from reactive, manual infrastructures into intent-driven, self-optimizing platforms. By utilizing intents, autonomous domains, artificial-intelligence-powered observability and robust conflict-resolution strategies, communication service providers can deliver reliable services efficiently while maintaining human oversight. As part of this evolution, the human operators' experience evolves from fragmented, task-oriented workflows to strategic stewardship, empowering them to guide outcomes rather than chase symptoms. Ultimately, AN operations will provide faster value delivery, resilient service quality and scalable operations, positioning telecom networks for the next wave of programmable innovation.

The authors

Ciaran Johnston is a senior expert in operations support systems (OSS) and programmable network architecture, and he is the chief architect of Ericsson's network management product portfolio. He joined Ericsson in 2000 and has over 20 years of experience in software development and architecture in the OSS domain. Johnston holds a B.Sc. in pure and applied physics from the University of Manchester in the UK.

Jörg Niemöller is an analytics and customer experience expert in the cognitive network solutions area. He joined Ericsson in 1998. He is currently leading the introduction of machine-reasoning technologies into Ericsson's portfolio to enable solutions for ANs. Niemöller holds a Ph.D. in computer science from Tilburg University, the Netherlands, and a Dipl.Ing. in electrical engineering from TU Dortmund University, Germany.

Ann-Christine Eriksson is an expert in RAN management systems and architecture. She joined Ericsson in 1988 and has worked in a variety of RAN-related research and development roles. She is currently responsible for technical strategies and the introduction of new technologies in the areas of RAN management and automation. Eriksson holds an M.Sc. in engineering physics and applied mathematics from KTH Royal Institute of Technology in Stockholm, Sweden.

Wenfeng Hu is a senior specialist in the cognitive network solutions area. He joined Ericsson in 2011 and has worked in various roles in service delivery and product development. He currently leads technology strategy with a focus on network digital twins, AI and mathematical optimization for network management. Hu holds an M.Sc. in wireless communication from Beihang University in Beijing, China.

Joseph Grogan joined Ericsson in 2012 and currently serves as a UX architect in the network management architecture group. With 15 years of experience working in the user interface and UX fields, he is currently responsible for developing strategies and technical solutions ensuring network management products support their broad range of target users. Grogan holds a B.A. in interactive media and web authoring from Cardiff Metropolitan University in the UK.

P.V.K. Ravikumar is a principal developer in the business support systems (BSS) and OSS domains. He joined Ericsson in 2008 and has worked in several technology areas and system architectures in charging, billing, orchestration, assurance, core commerce, AI and cloud-native domains. In his current role, he leads a team of senior architects building long-term BSS/OSS target architecture and evolution. Ravikumar holds a B. Tech. in mechanical engineering from Sri Venkateswara Hindu College of Engineering in Machilipatnam, India, and has completed a product manager certification program at Indian School of Business in Hyderabad, India.

References

1. Autonomy by design – enabling self-managing networks across the life cycle ↗
2. Ericsson Technology Review, Creating autonomous networks with intent-based closed loops, April 19, 2022, Niemöller, J.; Szabó, R.; Zahemszky, A.; Roeland, D. ↗
3. Ericsson white paper, Intent-driven networks is a key step in the journey to autonomous networks, February 2025 ↗
4. TM Forum, TR292 TM Forum Intent Ontology (TIO) v3.6.0 – TM Forum, August 2024 ↗
5. Ericsson white paper, AI agents in the telecommunication network architecture, October 2025 ↗
6. Ericsson Technology Review, Data ingestion architecture for telecom applications, March 16, 2021 ↗
7. Ericsson Technology Review, rApps: Transforming network management with intelligent automation apps, December 6, 2023 ↗
8. 3GPP TS 28.312, Specification # 28.312 ↗

Further reading

- Autonomous networks ↗
- Automation and AI ↗
- Talk to your network: Igniting the AI revolution in autonomous networks ↗