Ericsson White Paper
GFTL-26:001117
July 2026

Networks that reason and act: Adapting Agentic AI for intent-driven, closed-loop autonomy

Content

Introduction

The telecommunications industry stands at an inflection point. Networks spanning 5G, edge computing, Internet of Things, and the early foundations of 6G are becoming increasingly complex and expensive. Individual tasks can be optimized through manual intervention, reactive fault resolution, rule-based workflows, and isolated artificial intelligence (AI) models. Still, the industry struggles to manage end-to-end network behavior, resolve cross-domain dependencies, and adapt to changing conditions. A more intelligent and coordinated operating model is required to meet the demands of emerging applications such as real-time adaptability, ultra-reliability, and personalized service delivery at scale.

This is accelerating the industry's shift toward autonomous networks: intent-driven, self-optimizing, and increasingly self-managing networks capable of translating business objectives into operational actions with limited human intervention. Progressing toward higher levels of autonomy requires an intelligent control layer that can reason across domains, coordinate actions, learn from outcomes, and operate safely within clearly defined boundaries.

Agentic AI is a key enabler of this transition. Unlike traditional AI systems that operate within narrow, predefined workflows, agentic AI systems are goal-directed. They can interpret intent, reason over complex network contexts, plan and execute multi-step actions, interact with tools and systems, and continuously adjust based on feedback. Applied to assurance, incident resolution, provisioning, and network optimization, AI agents can move operations from reactive automation toward proactive, closed-loop management. They can reduce operating costs, shorten mean time to resolution, improve service quality and resilience, and create a foundation for more differentiated and outcome-based services.

However, telecom networks cannot simply adopt generic AI agents. Networks are critical infrastructure, highly regulated, and operationally sensitive. Agents must understand telecom-specific contexts, interact reliably with legacy and cloud-native systems, coordinate across multi-vendor and multi-domain environments, and deliver predictable outcomes without compromising security, resilience, or accountability. Human oversight will remain essential, particularly as operators progress from controlled use cases toward higher levels of autonomy.

This white paper explores how agentic AI can be adapted to accelerate the evolution toward autonomous networks. It outlines the operating principles and safeguards required to make AI agents suited for telecom networks: safe, auditable, interoperable, and scalable. It also provides practical design principles for moving from pilots to trusted production deployments without compromising network reliability or trust.

Agentic AI value to network operators

Agentic AI expands communications service providers (CSPs) value beyond cost optimization, unlocking platform monetization and new AI-native services, enabling both higher efficiency and new revenue streams. The largest upside lies in enterprise services, network APIs, and autonomous offerings that go beyond traditional connectivity revenue.

New business models

Agentic AI creates opportunities for CSPs to package network capabilities as outcome-based, AI-native services rather than static connectivity products. Examples include autonomous enterprise connectivity, self-optimizing private networks, service-level agreement (SLA)-backed application performance, security-aware connectivity, and managed digital operations for vertical industries. In these models, CSPs can monetize not only bandwidth but also assurance, automation, intent fulfillment, and guaranteed business outcomes. This opens the door to recurring revenue models, usage-based charging, premium SLA tiers, and managed agentic services for enterprises.

Network efficiency and cost reduction

Agentic AI could materially reduce operational expenditure by automating complex network workflows such as zero-touch provisioning, SLA-driven assurance, and autonomous fault resolution. Agents can continuously monitor network state, detect degradation before it affects customers, and execute remediation without human intervention. This would reduce mean time to resolution, decrease outage frequency and duration, resulting in higher network uptime, a better user experience, and a leaner, more resilient operation that scales without proportional growth in headcount.

Hyper-personalized service delivery

Agentic AI enables a fundamentally new approach to customer service. Specialized agents automate decision-making across multiple workflows, making personalization a possibility at the individual level and at scale. Rather than broad, undifferentiated campaigns, operators can target customers based on real-time context, behavioral signals, and inferred needs, delivering differentiated connectivity experiences that command higher margins and while improving retention.

Platform monetization

Network APIs can be exposed through traditional APIs or through agent-friendly interfaces such as Model Context Protocol (MCP) or Agent-to-Agent (A2A). As enterprise and developer workflows become increasingly AI-assisted, API discoverability, orchestration, and ecosystem richness will matter as much as the interfaces themselves.

Over time, these agent-friendly interfaces can evolve into teleco-grade agentic AI platforms built on top of existing networks. Such platforms would integrate network data, exposure APIs, AI services, policy controls, and network automation interfaces into a programmable environment for enterprises, developers, application providers, and partners.

These platforms can also enable privacy-preserving data and insight monetization. Rather than exposing raw data, CSPs can provide aggregated insights, predictive analytics, and agent-enabled decision support based on network, mobility, quality-of-experience, and service-usage signals. This expands the addressable market for CSP network capabilities and positions CSPs as an active participant in digital value creation.

Challenges in deploying agentic AI in telecom

To realize value of agentic AI for CSPs, several key challenges must be addressed to enable the effective integration of AI agents in telecom environments. This section provides a high-level overview of the key challenges associated with telecom agentic AI.

Safety and reliability: Telecommunications infrastructure is critical. Autonomous agents must operate within strict safety boundaries, with robust guardrails that prevent actions capable of causing service outages or regulatory violations.

Real-time constraints: Many network operations require decision-making on millisecond-to-second timescales, demanding low-latency inference and minimal coordination overhead amongst agents.

Multi-vendor complexity: CSP networks are inherently multi-vendor environments. Agentic AI systems must interoperate across diverse equipment, interfaces, and management systems without assuming a homogeneous stack.

Data sovereignty and privacy: Network data is highly sensitive. Agentic architecture must address data residency requirements, privacy regulations such as the General Data Protection Regulation, and the need to derive intelligence without exposing raw customer data.

Infrastructure requirements: Unlike stateless inference, agentic AI requires stateful, long-running workflows with continuous perception, reasoning, and multi-step tool coordination. Bottlenecks shift from pure compute to system-level orchestration, memory and context management, data locality, and low-latency agent communication. Telecom environments additionally require deterministic performance, edge-to-core orchestration, and telecom-grade governance covering SLAs, safety, and data sovereignty.

Legacy integration: CSPs operate systems with decades-long lifecycles. Agentic AI must integrate with existing OSS/BSS stacks rather than replace them, requiring carefully designed bridge strategies.

Adapting agentic AI for telecom applications

Telecommunications networks follow standardized functional architectures and implement standardized interfaces to enable interoperability in highly heterogeneous, multi-vendor deployments. AI agents must operate within this environment, implementing defined functions in network operations and participating autonomously in operational processes using the full range of established APIs, data sources, and management systems. This is the starting point of an evolution—from today's partially automated networks to highly adaptive, autonomous networks with AI at their core, realizing gains in cost efficiency, energy efficiency, and resource utilization.

Agent roles and scope

AI agents are a new class of components that must fit within established network operations architectures. Their primary strength is situational awareness and decision-making: analyzing situations, detecting opportunities to improve network state, and generating action plans. Their ability to handle complex, multi-variable scenarios can outperform traditional policy-based automation and, in many cases, human operators working under pressure.

They can take on central roles in control loops, observing network state, synthesizing action plans, and making decisions that balance intent, business needs, and technical constraints. Agent fabrics expose tools, data, and enablers to AI agents, while agents implement coordination, analysis, and decision-making. The scope of an agent fabric aligns with established operational domains—RAN, core, transport, cloud management, and end-to-end service management—preserving the modularity that makes large-scale telecom operations manageable. Each agent fabric exposes domain-specific models, tools, and data in support of a distinct group of agents that is operating within the domain. Furthermore, not all functions are best performed by agents; processes with strict real-time requirements, established control systems with static goals, and data-processing pipelines are better handled by dedicated systems that serve as enablers to the agentic intelligence layer.

Agent discovery and interaction

The telecom industry can leverage existing agent interfaces such as A2A for agentic communication. These are lightweight protocols primarily concerned with task-level conversations that allow one agent to utilize another's capabilities. These application programming interfaces (APIs) must remain lean and compatible without being overloaded with industry-specific features. Gaps such as agent authorization and capability alignment are best addressed through dedicated complementary interfaces for authentication, authorization, and language negotiation.

For correct interaction, agents must agree not only on vocabulary but also on semantics since agents on both sides must derive the same meaning from a conversation to take appropriate actions. Verifying semantic alignment is a new challenge in system integration that requires formally defined, constrained domain languages rather than open-ended natural language exchanges.

Designing safe agentic systems

Telecommunications networks are critical infrastructure, imposing high demands for consistency and correctness, while generative AI is non-deterministic and prone to hallucinations. Its flexibility makes agents powerful but also introduces security concerns. Agents must be subject to strong authentication mechanisms and clearly defined roles that determine the tasks they may request and the type of information they may share or receive.

Enforcing these boundaries is challenging because of the flexibility of the language used in agentic communication. Authorization must hold even when agents rephrase requests to find alternative paths to restricted information or capabilities. This represents a paradigm shift in security, from API-level access control toward content-level policing of agent conversations. Keeping the language constrained and formally defined enables reliable enforcement and auditability.

Keeping the network consistent

Zero-trust principles require not only that participating agents are authenticated and authorized, but also that their actions are independently verified as beneficial before execution. This counteracts both hallucination-driven errors and malicious actions resulting from external attacks such as prompt injection. An agentic environment asks not only what an agent wants to do, but also why and what outcome is expected, enabling independent verification before any action is taken.

Network digital twins are a key enabler, allowing proposed actions to be assessed in a safe sandbox before production execution. Intent-based requirements, including business utility specifications, provide high-level guardrails that keep agents within beneficial operational boundaries, complementing action-level guardrails that directly constrain agent actions to safe operating ranges.

Shifting human roles and trust

Autonomous networks enable greater operational efficiency by minimizing the need for direct human intervention. To achieve their full potential, organizations must be confident about autonomous decision-making while ensuring business objectives such as scalability, faster delivery, and cost efficiency are met. The following three capabilities are the key trust drivers that facilitate this evolution:

Policies that define goals and boundaries: Humans control what autonomous agents pursue by defining operational goals, business requirements, and priorities through intent. Policy guardrails keep operations within known safe boundaries, giving operators confidence without requiring step-by-step supervision.

AI explainability and self-awareness: Continuous monitoring and verification of operational decisions require AI explainability techniques combined with natural language interaction. Agents must proactively reach out to humans when they detect situations that are beyond their capabilities, rather than hallucinating inaccurate solutions. This makes agents transparent, reliable partners to the humans who remain ultimately responsible.

AgentOps with governance and lifecycle management: Safe production use requires carefully designed governance and lifecycle management processes, defining agent roles and scopes, sourcing and designing agents and tools, and verifying interoperability and security before production integration.

Infrastructure requirements

Agents based on generative AI require a deployment environment capable of handling significant compute and storage workloads, with specialized hardware where necessary. Agentic AI needs access to rich, contextualized data across network domains, demanding investment in data pipelines, normalization, and real-time streaming infrastructure.

Unlike stateless inference, agentic AI requires stateful, long-running workflows with continuous perception, reasoning, and multi-step tool coordination. Bottlenecks shift from pure compute to system-level orchestration, memory and context management, data locality, and low-latency communication among agents, models, and network functions.

While AI agents act and react with low latency compared with human-driven processes, their inference and interaction latency can still be a limiting factor for real-time requirements at a millisecond scale. The choice and combination of model technologies and their inherent performance characteristics need to match the use cases.

Telecom environments must ensure deterministic performance, edge-to-core orchestration, and telecom-grade governance covering SLAs, safety, and data sovereignty.

Domain contextual knowledge

Deterministic decision-making in agents depends heavily on the contextual information available to them. Telecom ontology-assisted reasoning reduces hallucinations and inefficient reasoning loops by grounding AI outputs in a structured, domain-specific model of network entities, KPIs, and causal relationships. This constrains the system to valid concepts, improves accuracy and explainability, and guides inference along focused paths such as fault → impact → root cause. Domain knowledge can be provided as a knowledge retrieval-augmented generation used by a frontier model, as fine-tuning data for a foundation model, or as the basis for building telecom-specific foundation models. Agentic knowledge might be derived from data provided by existing interfaces in the telecom operations stack.

Enabling legacy systems

CSPs operate systems with decades-long lifecycles. Agentic AI must integrate with existing OSS/BSS stacks rather than replace them. The automation journey may begin with modern systems where acquiring information and changing behavior at runtime is straightforward, then progressively extend to legacy systems from multiple vendors built on older technologies. The following four approaches address this challenge:

- Expose legacy interfaces as tools: Define tool schemas clearly through standalone MCP server implementations that act as proxies, making legacy capabilities available to agents without modifying the underlying systems. However, MCP is not suitable for all layers of the telecom stack, and in some cases, it is preferable to expose functionality through APIs.
- AI agents that work with legacy systems: Enable AI agents to understand the scope and tasks managed by legacy systems. They can interact with legacy systems through their native APIs and scope their own operations accordingly. AI Agents work with and around legacy systems that need to be kept in place by the CSP.
- Create actuation hooks: Establish mechanisms for legacy systems to receive commands, through available interfaces or, where necessary, through command-line interface commands, enabling agents to effect change across the full network stack.
- Model legacy entities in the knowledge ontology: Ensure legacy systems are semantically connected to other applications through shared knowledge, so that agents do not need to treat them differently from modern systems.

Measuring agent value

Measuring the value of agentic systems requires metrics that capture both operational outcomes and business impact. Relevant indicators include reductions in mean time to resolution, percentage of incidents resolved autonomously, OPEX savings from zero-touch operations, revenue uplift from personalized service delivery, and developer productivity gains from autonomous tooling. Establishing baselines before deployment and continuously tracking these metrics enables CSPs to demonstrate return on investment and make informed decisions about where to expand agentic automation.

Demonstrating the value of agentic AI

Autonomous network optimization with agentic rApps

Network optimization is a multi-objective task. For example, user experience optimization focuses on providing consistent, and uninterrupted access to sufficient throughput and low latency, while energy optimization eliminates wasteful resource utilization.

Ericsson realizes multiple optimization concerns through dedicated optimizer rApps following the service management and orchestration (SMO) architecture specified by open RAN. It provides a flexible environment for deployment and lifecycle management of autonomous software components. The rApps implement optimization objectives, including root-cause analysis and impact evaluation. See Figure 1.

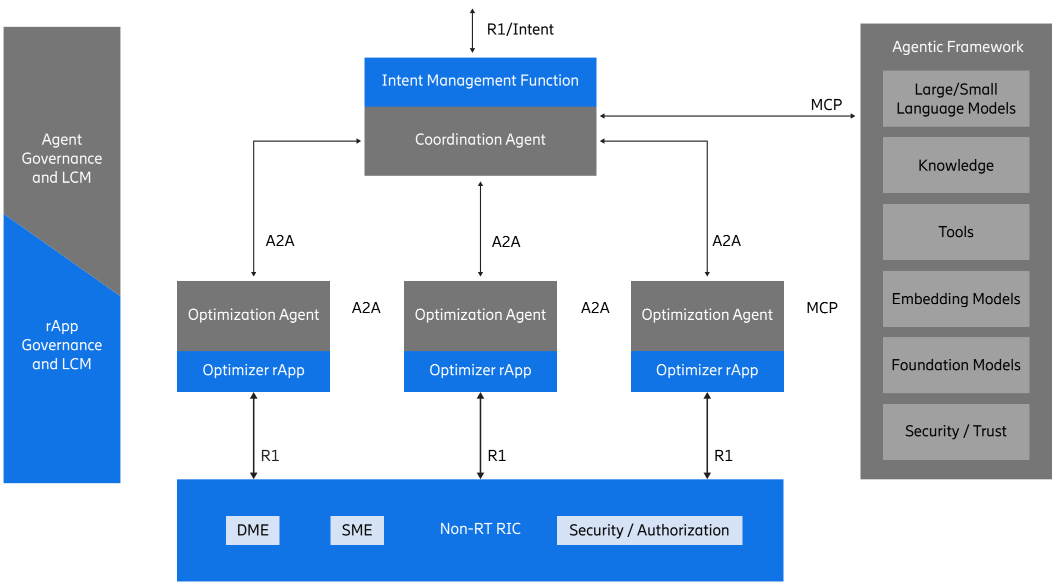

Figure 1. Orchestration of Agentic rApps

Network optimization is a multi-objective challenge and contributing optimization rApps need coordination toward an overall preferred network state. Ericsson introduces agentic AI capabilities into rApps and adds a coordination agent. This allows rApps to interact as AI agents with each other, the coordinator and an agentic framework through protocols such as A2A and MCP. The coordinator agent also includes an intent management function responsible for intent fulfilment, including collaborative conflict resolution and utility-driven solution finding. While the agentic AI overlay is concerned with collaboratively finding optimized solutions, agentic rApps remain fully integrated with the SMO and Non-Real-Time RAN Intelligent Controller through the standardized R1 interface.

Evolving rApps into agentic AI rApps demonstrates a gradual evolution strategy, raising overall performance and capabilities. Established life cycle management and governance of rApps directly applies to agents providing a deployment environment. Standardized SMOs facilities for providing data, intent, and management access to network resources enable agents to stay informed and act. On the other hand, rApps benefit from flexible and intelligent agentic communication when facing conflicts that require utility-driven coordination.

Autonomous network anomaly resolution

Traditional approaches to network anomaly resolution rely on predefined alarm-correlation rules and manual root-cause analysis—processes that are slow, error-prone, and unable to keep pace with network complexity. Agentic AI-based RAN troubleshooting orchestration continuously ingests and analyzes multi-domain network data, including performance counters, alarms, configuration changes, and environmental data. Domain-specific agents covering RAN, transport, core, edge, virtualization, and hardware maintain deep expertise in their respective domains. Agent and tool orchestration flows are driven by a knowledge graph derived from the domain and functional ontology.

In this automation use case, a detected anomaly is transformed into a system prompt that starts the resolution process. A reasoning agent uses knowledge graph to create a resolution plan and begins invoking other agents and tools to reason across network topology, temporal correlations, and historical incident patterns. After root cause analysis is completed, a remediation plan is generated with a list of actions. Some actions are executed by actuation agents, and some actions are escalated to human operators based on policy directives, which ensures that KPIs are not degraded by a proposed configuration or action recommendations. Agents and tools are adhering to standardized communication protocols that enable multi-party agent and tool interplay (Figure 2). Early deployments have demonstrated significant reductions in mean time to resolution and increases in first-contact resolution rates, with the potential to automate the resolution of up to 70 percent for routine network incidents.

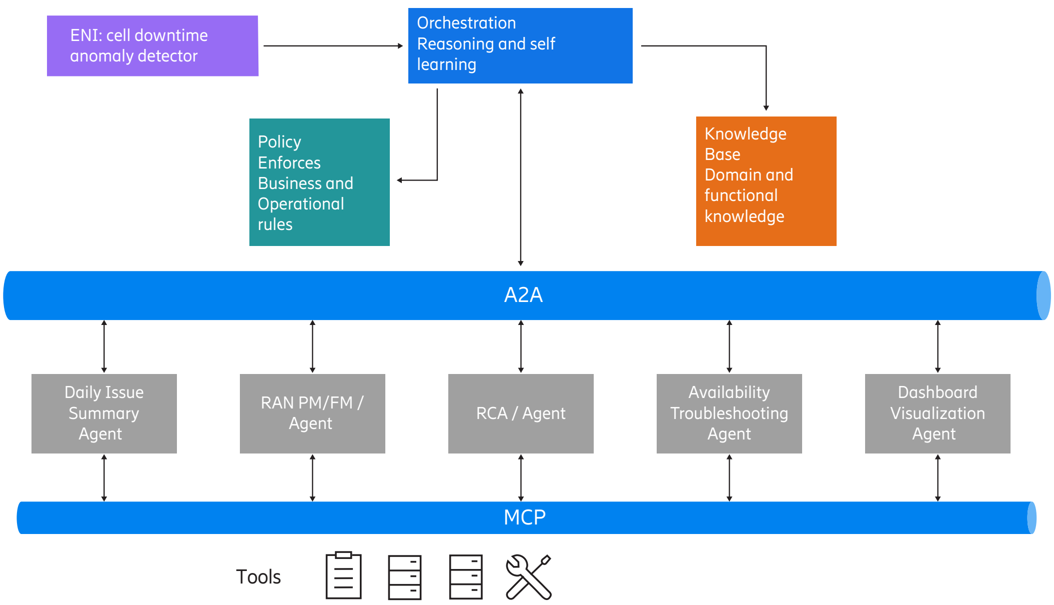

Figure 2. RAN troubleshooting with specialized agents and tools

End-to-end cross-domain agentic service intent management

Cross-domain service intent handling is a multi-faceted challenge that is difficult to realize using traditional rule-based or policy-driven frameworks. The process itself involves multi-stage analysis and identifying the most probable solution to address the observed service degradation. Each stage must work autonomously and determine at runtime, which is the next stage, who is best equipped to handle it, and how to delegate accordingly. Without agentic AI, the complete cognitive loop must be hardcoded with rules and leaving no scope for the process to adapt the loop for unanticipated network behavior or anomalies.

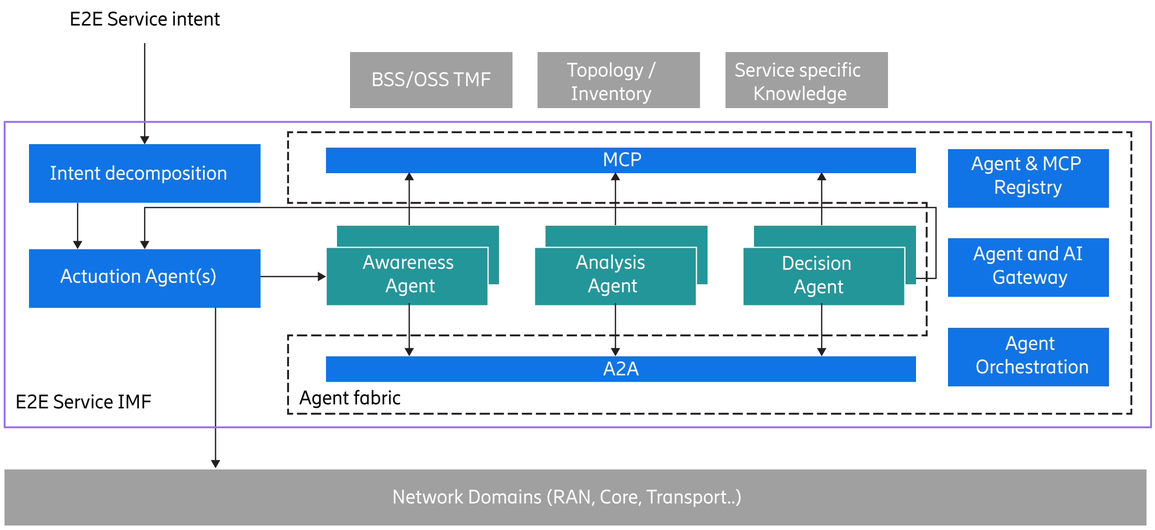

Figure 3. Agents for cross domain intent handling

In the context of service intent management, an agent—based on the outcomes it derived—discovers other agents through agent registry that can perform the next task based on their advertised capabilities (Figure 3). This is not a hard-coded sequence but an autonomous decision-making process in which agents dynamically hand over the task to one another. Agents exchange tasks through agentic protocols such as A2A. They are also assisted in decision-making process through the agent orchestrator in achieving the overall service-level assurance.

At each stage, these agents are also expected to make decisions within their scope based on their analysis and the same is not rule driven. They interact with required tools and resources through standardized mechanisms such as MCP and certainly hard wiring with tools would make the closed loop static. Agentic methods introduce autonomy and adaptive behaviors to complex functionality.

Conclusion and call for action

Agentic AI represents a major technological shift in telecommunications. It is the foundation for a new operating model where networks can reason, adapt, and act with increasing autonomy. As networks become more complex, dynamic, and service-critical, CSPs will need systems that can perceive, reason, plan, and act with increasing independence. The impact spans the full telecom value chain: lower costs, faster innovation, better customer experience, and new revenue opportunities.

As AI agents become capable of acting on critical systems, the value shifts from intelligence alone to trusted deployment, governance, and controlled autonomy.

CSPs must empower agents with the right tools, context, and autonomy to act at operational speed, while controlling them through strong governance, least-privilege access, policy enforcement, verification, and human escalation paths. With end-to-end observability, digital twin testing, continuous validation, and auditability, agentic AI can become a trusted control fabric for autonomous networks.

The industry must also avoid a fragmented future of proprietary agents, closed interfaces, and isolated automation islands. Scaling agentic AI across telecom will require open standards for agent interfaces, shared data semantics, governance, lifecycle management, and validation. CSPs and vendors that build on this foundation will be best positioned to lead the transition toward autonomous networks, unlock differentiated enterprise services, and capture the AI-driven value pools emerging around 5G and future 6G networks.

Beyond network operations, agentic AI is also redefining the enterprise API marketplace: the primary consumer is shifting from human developers to autonomous agents, consumption from hard-coded integrations to dynamic reasoning-based workflows, and value from data delivery to resolved business outcomes. This elevates the importance of trusted context, orchestration, and execution environments and creates a new opportunity for telecom players to extend their role from connectivity providers to outcome-enabling platforms.

References and further reading

1. Whitepaper on AI agents in the telecommunication network architecture.
2. Whitepaper on Trustworthy AI for Telecom systems.
3. Four ways agentic AI will reshape the telecom ecosystem.
4. Agent2Agent (A2A) Protocol: A2A protocol specification v1.0.0.
5. Model Context Protocol (MCP), specification and documentation for the model context protocol.
6. Intent-driven networks is a key step in the journey to autonomous networks

Authors

Ursula Challita is an AI Strategy Manager in Ericsson’s Group Function Technology Leadership organization. She joined Ericsson in 2018 as a researcher focused on AI for RAN optimization and has since held several roles including AI technical lead and AI strategi driver across multiple portfolio offerings. She holds a PhD from The University of Edinburgh, UK, in AI for Telecom networks.

Jörg Niemöller is an expert in analytics and customer experience. Since 1998, Jörg has held positions in Ericsson Research, core network system management, and digital services. He has developed concepts and solutions for intelligent systems capable of driving autonomous operation and realizing the zero-touch vision. His current focus is the introduction of these technologies in our industry through evolved products and standardization. Jörg is the lead author of TM Forum guidebooks and models on intent.

Zeljka Lemaster is the Head of the AI Innovation and Incubation US Hub at Ericsson. She leads strategic initiatives focused on applying latest AI technologies to telecom networks and operations. Her team's most recent work spans the development of AI-native frameworks, intelligent automation solutions, and autonomous network capabilities that leverage foundation models, domain-specific knowledge, and multi-agent systems to accelerate innovation across Ericsson's portfolio and deliver measurable value to communication service providers. Zeljka also leads AI ecosystem engagement efforts across North America and Silicon Valley, collaborating with technology leaders, startups, hyperscalers, research institutions, and open-source communities.

P.V.K. Ravikumar is a Senior Expert of Telecom systems architecture with more than 25 years of experience spanning across monetization, orchestration, assurance, and core commerce domains within BSS/OSS. Ravi is presently leading a group of senior architects to define and drive future architecture direction of Ericsson BSS and OSS application suite considering various industry, technology, and market trends. He is responsible to establish OSS and BSS architecture evolutions in the larger Ericsson telecom portfolio landscape. Ravi's educational background includes Bachelor of Engineering.