NGMN

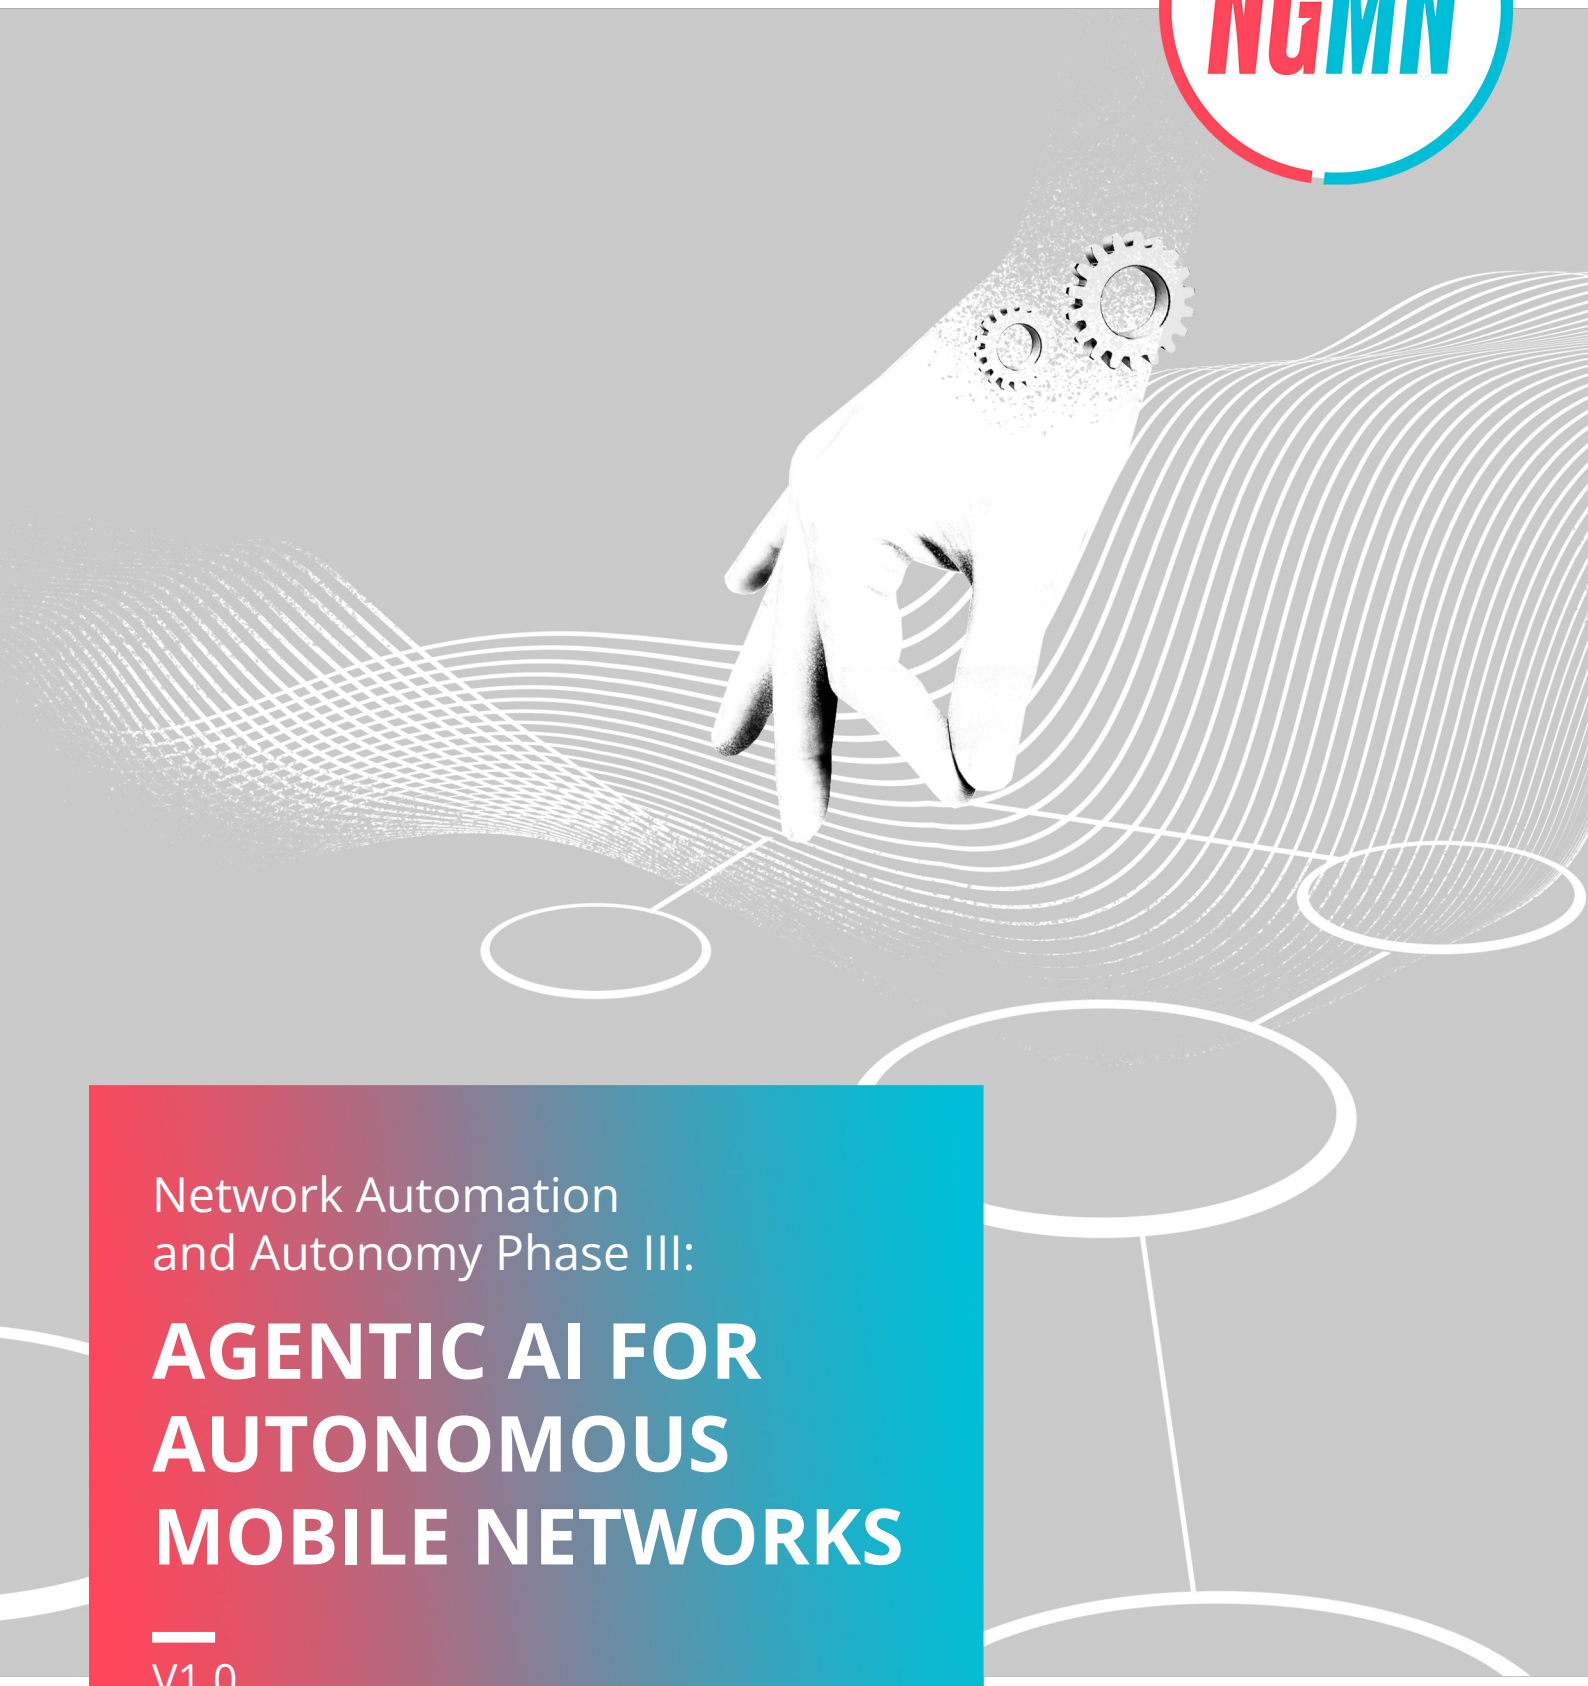

Network Automation
and Autonomy Phase III:

AGENTIC AI FOR AUTONOMOUS MOBILE NETWORKS

V1.0

ngmn.org

WE MAKE BETTER CONNECTIONS

Network Automation and Autonomy Phase III:

AGENTIC AI FOR AUTONOMOUS MOBILE NETWORKS

by NGMN Alliance

Version:	1.0
Date:	11 August 2026
Document Type:	Public
Programme:	Mastering the Route to Disaggregation
Approved by / Date:	NGMN Board, 5 August 2026

CONTENTS

EXECUTIVE SUMMARY 4

01 INTRODUCTION 5

1.1 Overview 5

1.2 Core Concepts: Agentic AI, AI Agents and Agentic Capabilities 5

1.3 Document Structure and Analytical Framework 6

1.4 Target Audience 7

02 PROBLEM STATEMENT AND BUSINESS DRIVERS 8

03 SELECTED SCENARIOS DEFINITION AND ANALYSIS 9

3.1 Fault Management 9

3.2 Service Assurance 11

3.3 RAN Optimisation 13

04 REFERENCE FRAMEWORK FOR AGENTIC AI COLLABORATION 15

4.1 Core Components 15

4.2 Agent Collaboration Patterns 17

4.3 Interface Requirements 17

4.4 Security and Trust Mechanisms 18

4.5 Alignment and Gap Analysis with Existing Standards 21

05 CONSIDERATIONS FOR THE ECOSYSTEM 23

5.1 Purpose and Positioning of Considerations 23

5.2 Considerations for Open-source Communities, Standards and Industry Organisations 23

5.3 Considerations for MNOs 24

5.4 Considerations for Vendors, Solution Providers and Hyperscalers 24

06 CONCLUSION 25

APPENDIX A: GLOSSARY 27

APPENDIX B: EXISTING STANDARDS, OPEN SROUCE PROJECTS AND INDUSTRY INITIATIVES 30

REFERENCES 35

FIGURES 40

ACKNOWLEDGEMENTS 41

EXECUTIVE SUMMARY

Mobile network operators (MNOs) now face a decisive network automation inflection point. Rule-based automation and machine learning (ML) tools have improved efficiency and enabled networks to become partially or conditionally autonomous, characterised by automated tasks in isolated domains, AI/ML assisting human decision making and humans in the loop at least for approvals. However, the current situation remains insufficient to effectively and efficiently automate dynamic, cross-domain high-value use cases, given the growing complexity of mobile networks and high expectations regarding customer experience. The next major target in our industry is Level 4 autonomy as per TM Forum [7] where networks are highly autonomous through proactive, within-domain and cross-domain intent-driven operations requiring minimal human intervention while enabling automated decision making, self-optimisation, self-healing and self-management.

This vision appears increasingly achievable through the rapid progress being made with generative AI, the growing capabilities of large and small language models (LLMs, SLMs), and the increasing effectiveness of collaborating AI agents. They interpret intent, reason over network and service context, plan, delegate tasks across agents, invoke existing tools and controllers, validate actions, coordinate recovery or optimisation across network domains and learn and adapt. Early forms of such multi-agent systems are already feasible today and are increasingly supported by a rapidly evolving ecosystem of AI agent platforms, frameworks and infrastructure.

While this NGMN publication positions Agentic AI as a foundational enabler of network automation in autonomous mobile networks, it highlights that achieving this vision requires much more than the technologies currently attracting the greatest industry attention, such as LLMs and AI agents.

First, Agentic AI is expected to complement and interact with rather than replace existing automation assets. Second, without reliable, comprehensive, complementary software support systems that enable governance, trust, policy adherence, observability, explainability, operational safety, determinism, assurance, security and cost control, the vision is neither achievable nor sustainable for network operators. In addition, automation with Agentic AI is only feasible at scale if interoperability is supported by suitable standards for agentic communication.

In this publication, the vision is motivated and then exemplified using three illustrative high-value scenarios. For decision makers, this complexity is reduced through an operator-centred reference framework that identifies the core capabilities, functions and enablers required for network automation with Agentic AI. From an exemplary gap analysis across the industry the publication derives important considerations for the whole ecosystem including standards organisations.

The strategic conclusion is clear: commercial-scale adoption will depend less on aspects like individual AI model capability than on ecosystem alignment, reduction of fragmentation risks and readiness of support systems for Agentic AI. To name a few, operators need interoperable agent communication, shared semantics, secure agent identity, auditable execution, digital-twin validation, lifecycle assurance, and practical human-in/on-the-loop controls. Standards development organisations, open-source communities, vendors, hyperscalers and operators should therefore converge among others on telecom-grade interfaces and trust mechanisms and close any potential or perceived gaps identified in this report. If these challenges are addressed, Agentic AI can reduce operational friction, improve service quality, support differentiated services and create a credible pathway from today's Level 2/3 automation as per TM Forum [7] toward trusted Level 4 and, eventually, Level 5 autonomous networks.

01 INTRODUCTION

1.1 OVERVIEW

This publication represents the output of an operator-driven initiative delivered by the Next Generation Mobile Networks Alliance in the first half of 2026. Its objective is to provide network operator-driven guidance to the global telecom ecosystem on deploying Agentic Artificial Intelligence to achieve advanced network automation within Autonomous Mobile Networks.

As mobile network complexity escalates, operators must rapidly evolve their operational paradigms. Today, most service providers operate at Autonomous Network (AN) Levels 2 or 3, as defined by TM Forum Autonomous Networks (AN) Framework [7]. All notations regarding automation levels in this publication refer to the TM Forum AN Framework. The next critical milestone is the transition to Level 4 autonomy. Network operators aspire to realise this shift over the next few years, individually targeting those high-value scenarios which they judge to yield the most compelling economic returns.

However, the path to implementation presents significant hurdles. Agentic AI technology is developing rapidly across a fragmented landscape of diverse organisations, while several critical requirements important to network operators appear to remain insufficiently addressed. These gaps must be resolved before operators can confidently commit to the internal transformations and projects required for Level 4 deployments in the years to come.

1.2 CORE CONCEPTS: AGENTIC AI, AI AGENTS AND AGENTIC CAPABILITIES

As network automation evolves from rule-based automation and AI-assisted tools towards Level 4 autonomous networks, it is important to clarify the relationship between AI Agents, Agentic AI and Agentic AI capabilities. In this publication, these terms are used in the context of mobile network operations and management, rather than as generic AI concepts. The purpose is to distinguish Agentic AI from conventional workflow automation, isolated AI/ML analytics, and closed-loop mechanisms that operate within predefined rules and limited domains.

An AI Agent is an autonomous or semi-autonomous software entity that can perceive operational inputs, reason over relevant context, make decisions or recommendations, perform or trigger actions to achieve defined objectives as well as learn and adapt over time. In mobile network operations, such inputs may include telemetry, alarms, events, tickets, intents, policies, service requirements, topology information, performance indicators, configuration parameters (and their allowed value ranges) and operational history. These inputs represent a solution space against which an AI agent can find the optimal target configuration and keep it up-to-date by tracking the evolving context. An AI Agent may interact with other agents, OSS/EMS functions, management services, digital twins, data platforms, automation scripts or domain controllers through controlled interfaces.

AI Agents often receive long-lived, business-oriented intents and translate those into situation-dependent, short-live goals. A user-level or system-level intent is an abstract declarative statement (e.g. maintain 15 ms latency), while a goal is an agent-internal objective, that can be imperative and short-lived as the exact goal (how to satisfy the intent) may change as network circumstances or context vary.

LLMs and SLMs or foundation models are expected to play a central enabling role in many high-value AI Agent and Agentic AI scenarios. They provide flexible reasoning, natural-language and intent understanding, planning support, explanation generation and cross-domain knowledge synthesis that are difficult to achieve through rule-based automation, fixed workflows or isolated domain-specific AI/ML models alone. In this sense, large models are a key source of the distinctive value of Agentic AI, especially when operational objectives are ambiguous, complex in-domain, cross-domain, multi-step or dependent on changing network and service context. At the same time, large models are expected to be combined with domain-specific AI/ML, knowledge graphs, optimisation algorithms, digital twins, deterministic controllers and existing

automation systems, particularly where telecom-grade latency, stability, safety and verifiability are required.

Agentic AI refers to a system paradigm in which one or more AI Agents pursue operational objectives with a degree of autonomy under operator-defined policies, guardrails and governance constraints. Compared with a single AI Agent designed for a bounded task, Agentic AI systems are typically characterised by specialised agents, task decomposition, shared or exchanged context, agent-to-agent collaboration, tool and management capability invocation, execution result, feedback from consumer agents to producer agents, and conflict coordination. In mobile network operations, this may involve, for example, a cross-domain supervisory agent coordinating RAN, core, transport, cloud, security or field-operation agents to diagnose service degradation, evaluate corrective actions, resolve conflicts and verify service recovery. Within a single domain like RAN, a supervisory agent may orchestrate multiple specialist types of RAN agents.

Agentic AI capability refers to the system-level operational capability created when AI Agents, large models, tools, data, management services, interfaces and governance mechanisms work together to support controlled autonomous operation. Such capability includes the ability to interpret operational objectives, reason over dynamic network and service context, plan and decompose tasks, discover and invoke appropriate tools or management capabilities, collaborate with other agents or systems, monitor execution, adapt based on feedback and execution result, and operate within defined policy, safety, security, trust and governance constraints.

Agentic AI is therefore expected to complement rather than replace existing network automation assets. Closed-control loops, intent handlers, analytics functions, digital twins, domain controllers, OSS workflows, RAN applications and management services may become tools, resources or controlled execution environments that agents can discover, select and invoke. The role of standards and open-source alignment is to make these capabilities interoperable, secure, observable, auditable and governable across vendors, domains and deployment environments.

1.3 DOCUMENT STRUCTURE AND ANALYTICAL FRAMEWORK

To guide the industry through this complex transition, this report is structured as follows:

- **Chapter 2 (Problem Statement and Business Drivers):**

  Formalises the technical challenges and establishes the imperatives compelling operators to adopt Agentic AI.

- **Chapter 3 (Selected Scenarios Definition and Analysis):**

  Exemplifies for three illustrative high-value scenarios their business value, operational pain points and limitations of existing automation, and outlines a vision of Level 4 autonomy driven by agentic collaboration.

- **Chapter 4 (Reference Framework for Agentic AI Collaboration):**

  Introduces a structured reference framework considered relevant for network automation. This framework deliberately occupies a pragmatic middle ground, providing actionable technical utility without overwhelming the reader with implementation-specific minutiae. Section 4.2 calls out AI Agent collaboration patterns, Section 4.3 summarises interface requirements and Section 4.4 addresses vital security and trust mechanisms. Section 4.5 on ecosystem landscape and gap analysis provides a brief review of the standardisation and open-source projects intersecting Agentic AI, Autonomous Networks, and Network Automation. This section reviews a complex landscape spanning, among others, TM Forum, ETSI, 3GPP, IETF, and Linux Foundation projects. Utilising the reference framework, it synthesises the network operator perspective to expose areas that may benefit from further industry collaboration and opportunities for additional development.

- **Chapter 5 (Considerations for the Ecosystem):**

  Leverages NGMN's positioning as a non-standards-defining organisation to deliver objective, actionable considerations to operators, vendors, SDOs and open-source communities.

- **Chapter 6 and Appendices (Conclusions, Supplementary Data):**

  Concludes the report and provides annexes for specialist readers.

1.4 TARGET AUDIENCE

This publication is designed for a diverse, multi-tiered stakeholder base across the telecommunications value chain:

* **CTOs and technical strategists:**

  to assist with corporate roadmaps, investment strategies, and architectural visions and blueprints.

* **Technical management and expert communities:**

  to assist in planning, scoping and roadmapping network automation projects within operators and vendors alike.

* **Standards developers and open-source contributors:**

  to raise awareness of relevant work happening across multiple organisations, to point to areas that require more work, and to ensure that emerging standards and open-source code bases are robust, interoperable, secure, reflect a wide consensus and are fully fit for purpose for MNOs.

This publication reflects operator perspectives intended to guide ecosystem-wide alignment. Ultimately, it shall contribute to building the foundational confidence and trust required for MNOs to adopt these technologies at commercial scale.

02 PROBLEM STATEMENT AND BUSINESS DRIVERS

Traditional network automation has improved the efficiency of network operations, but it is generally based on predefined rules, fixed workflows and explicit triggering conditions. This approach is effective for repetitive tasks with well-defined boundaries, yet it is less able to adapt to dynamic service requirements and changing network conditions. When an issue falls outside the scope of existing rules, operations personnel may need to reassess the context, modify the workflow or develop additional scripts. Most automation capabilities have also been developed for individual network domains or management systems, with limited coordination between local closed loops. Consequently, even where individual steps are automated, end-to-end operational processes may still depend on manual analysis and cross-team coordination. This increases handling time and operational cost and constrains the evolution towards higher autonomous network levels.

Differentiated services further increase this complexity. Operational objectives are shifting from maintaining generic network performance towards assuring end-to-end outcomes for specific service requirements. Different services may place different emphasis on user experience, reliability, latency or cost, and their priorities may change with customer requirements and network conditions. The fulfilment of domain-specific KPIs therefore does not necessarily mean that end-to-end service requirements are met. An optimisation action within one domain may also affect another domain, the overall service or the ability to achieve a cross-domain operational objective. Network operators need to relate service objectives more directly to network conditions and the risks associated with corrective actions, and coordinate the required capabilities according to the service impact. Fixed workflows are not sufficient for such dynamic decision processes. Network operations need to evolve from local automation towards end-to-end coordination and controlled closed-loop automation driven by desired service outcomes.

Agentic AI provides a new approach to addressing this complexity and offers strategic value in three areas. First, it can organise complex operational activities around a common objective, enabling specialised agents to divide responsibilities, invoke existing management and automation capabilities, and coordinate actions within a controlled workflow in a multi-agent environment. An operational objective can be decomposed into assignable and traceable tasks, while each agent selects the appropriate capability according to the task context. Capability descriptions such as an Agent Card can declare an agent's role, skills, supported tasks and applicable conditions, thereby supporting skills discovery and matching of skills to tasks. This can reduce manual investigation, repeated operations and cross-team communication, shorten fault handling and service restoration time, and reduce operational cost.

At the same time, earlier detection, root-cause analysis and coordinated recovery can reduce the impact of service degradation on users, thereby improving network quality, service continuity and customer experience. Second, Agentic AI can dynamically organise operational capabilities according to service requirements, allowing service assurance to respond more precisely to specific customers and scenarios rather than relying only on common rules. This provides a more flexible basis for service differentiation. Third, network operators may progressively expose network capabilities, operational knowledge and agent capabilities as reusable services. Establishing a governable and scalable agent collaboration environment can strengthen the operator's technical influence and support a leading role in the evolution of autonomous networks and the broader industry ecosystem.

However, large-scale adoption of Agentic AI is constrained by results accuracy and thus security or trustworthiness concerns as well as by fragmented standards, open-source initiatives and proprietary solutions, which currently provide no consistent framework for mobile network operations. Without alignment on capability descriptions, task models and collaboration mechanisms, multi-vendor environments risk becoming isolated agent ecosystems with limited interoperability and governance. Addressing these gaps is the primary business driver for the project's scenario analysis, framework development and gap assessment.

03 SELECTED SCENARIOS DEFINITION AND ANALYSIS

This chapter introduces three illustrative high-value scenarios in the context of network automation and sketches a possible target state for L4 network autonomy for each of them.

TM Forum has collected and documented numerous high-value use cases and scenarios in its IG1339 Autonomous Networks L4 High-Value Scenarios guide [1]. For the purpose of illustrating the application of Agentic AI in network automation, this document focuses on three example scenarios: Fault Management, Service Assurance, and RAN Optimisation. The selection of these three scenarios should not be interpreted as an NGMN assessment of their relative priority or importance. Such priorities will vary across MNOs according to their individual business objectives, operational environments, and levels of automation maturity.

Importantly, this scenario-focused chapter does not discuss underlying network and IT architectures. Instead, the scenarios are presented through exemplary high-level workflows. These workflows are not intended to be comprehensive or to cover every operational situation that may arise in practice. Rather, their purpose is to demonstrate how key elements of Agentic AI, such as different types of agents, agent hierarchies, tools, and digital twins, could be applied within representative workflows for each scenario. In addition, the examples highlight potential collaboration patterns between agents. For each scenario we motivate its business value, highlight current limitations, propose a possible target state for L4 network autonomy, and identify relevant agent collaboration patterns and key technical enablers.

This chapter therefore serves primarily as an introduction to the reference framework presented later in this document. The framework subsequently elaborates on the key capabilities, functions, and components that are likely to be required to enable network automation through Agentic AI.

3.1 FAULT MANAGEMENT

While a more forward-looking approach may fuse fault management and performance management into network observability and integrate it with service assurance to better understand service impact, we limit the illustration here to fault management for the sake of simplicity.

Business value:

In the evolution of Autonomous Networks (AN) toward L4 high-level autonomy, fault management - in-domain/single-domain and especially cross-domain complex fault handling - is a high-value scenario for operators, directly impacting customer experience, network reliability, MTTR, and operational costs.

Current state and limitations:

Current core network fault management has generally reached L3, achieving rule-based automation but facing limitations: preset rules struggle with dynamic changes, automation is siloed, and cross-domain issues require extensive manual coordination, creating an efficiency bottleneck that can lead to lengthy manual resolution. Human oversight required for major decisions adds delay. These constraints result in slower responses, higher operations costs, and sometimes prolonged service impacts.

Level 4 target state:

To break through the L3 bottleneck and achieve a qualitative leap to L4, the industry is focusing on a new architecture and process reconstruction centred on Agentic AI. The core goal of the L4 stage is to achieve end-to-end full-process automatic closure, autonomous handling of routine and complex faults within human-set policy bounds, only involving humans for exceptions, deep autonomous collaboration of group intelligence, and precise virtual control through digital twins, driving fault management from a "ticket-driven, human-centric" model towards an "intent-driven, ticket-less" autonomous operational stage. Guided by this

goal, the cross-domain fault self-healing scenario constructs a collaborative system consisting of a cross-domain collaboration agent (or supervisory agent) and various professional domain (wireless, transmission, core network) fault handling agents, able to perform fault diagnosis and remediation, and embedded with large model analysis and reasoning capabilities. This scenario is complemented with AI co-pilots for occasions when human engineers must be involved (e.g. for on-site repairs) to accelerate the remaining manual resolution tasks. Depending on the conditions triggering the self-healing process, this scenario includes two typical and complementary fault self-healing sub-processes, each reflecting different emphases of the L4 goals: "proactive prevention" and "collaborative closure":

- Sub-scenario 1:

  OSS-triggered self-healing. This embodies end-to-end automatic closure and a customer-centric approach. When the OSS detects a service interruption, the cross-domain collaboration agent (central orchestrator or supervisory agent) initiates secure interaction with domain agents, dispatches investigation tasks, and coordinates analysis. Domain agents analyse their state, generate repair solutions, and execute them based on risk assessment (with automatic execution or human-in/on-the-loop for high-risk steps). Finally, results are aggregated, and service recovery is verified. This transforms fault handling from passive response to an active, business-outcome-driven AI process.

- Sub-scenario 2:

  Domain-controller- or EMS-triggered self-healing. This highlights deep group intelligence collaboration and endogenous network intelligence. When domain agents at the domain controller or EMS level detect but cannot resolve a fault locally, they securely report to the cross-domain collaboration agent. The cross-domain collaboration agent performs cross-domain correlation analysis, generates a collaborative repair solution, and coordinates execution (via other agents or field dispatch). This enables early detection and multi-domain collaborative decision-making.

In the example of sub-scenario 1 shown in Figure 1, the workflow is triggered by a customer incident ticket (1). A cross-domain fault management agent interprets the intent and delegates a task to a fault agent for the RAN domain (2). Within the RAN domain, the fault agent acts as a supervisor agent and delegates a task to a sub-agent (3). Inside the RAN autonomous domain, sub-agents may query and share information directly between them (4). Agents can also invoke tools and digital twins. The response from the RAN fault agent may indicate a root cause inside the RAN, however with only limited confidence. The cross-domain agent next invokes a Core network agent (5), which in turn delegates a task to a lower-level agent (6). If the RAN agent and the Core agent suggest different root causes and possibly conflicting remedial actions, the cross-domain agent performs correlation and drives any required conflict resolution (7). Throughout, MNOs are accountable for all actions initiated by agents. The issue resolution may require human action through field operations (8), where field engineers are assisted by AI copilots.

This example workflow also illustrates sub-scenario 2, should it commence at step (3), with the RAN fault agent autonomously detecting an incident and proceeding without a customer-triggered ticket.

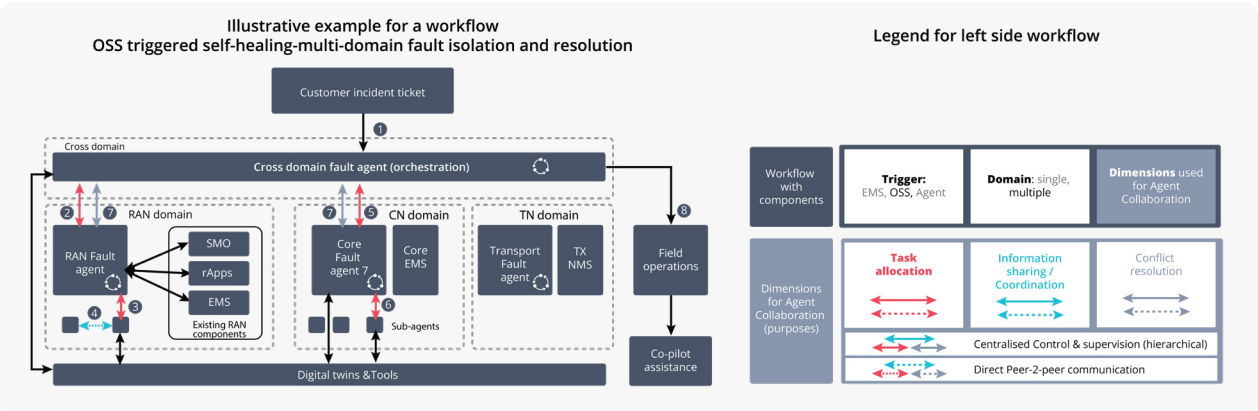

Figure 1: Illustrative example for a fault management workflow

Agent collaboration patterns:

From this perspective, the multi-agent interaction in this scenario follows a distinct model. In task allocation, it belongs to a centralised mode, whereby the OSS-based cross-domain collaboration agent adopts centralised management to assign investigation, collection, or repair tasks to domain-specific agents. For information sharing, the multi-agent interaction primarily employs a direct mode, where agents exchange fault reports, status data, and analysis results directly through standardised , secure channels (e.g., based on A2A-T interfaces [8]). In some cases, this may even enable peer-to-peer coordination. Publishing states and analyses into a shared dataspace and reading from this space is an alternative method. Regarding conflict coordination, in this example it also uses a centralised approach; the central collaboration agent arbitrates and makes decisions based on global policies when ambiguities arise in cross-domain analysis or when repair solutions involve resource or permission conflicts. This "task allocation–centralised , information sharing–direct, conflict coordination–centralised " pattern effectively underpins the end-to-end automation from fault perception to closed-loop recovery, representing a typical architectural embodiment of L4 deep swarm intelligence collaboration in fault management. Besides multi-agent interaction, agents routinely invoke external tools and services (e.g. a digital twin).

Key transformation enablers:

The gradual moves from L3 (or sublevels L3.x like 3.2, 3.4) to L4 are systemic, marked by progress across four interconnected dimensions. This transformation is driven by i) integrating large models with knowledge graphs to convert fragmented operational knowledge into a reasoning-capable system, enabling precise handling of complex and long-tail faults. ii) It relies on standardised agent interfaces, such as A2A [29] and A2A-T [8], to facilitate cross-layer and cross-domain task negotiation and dynamic orchestration, breaking down system silos and shifting from serialised manual handovers to parallelised intelligent collaboration. iii) Agents function as digital employees, dynamically converting textual procedures into executable sequences to achieve end-to-end automatic closure, thereby moving operations from a ticket-driven model to intent-driven, autonomous execution. iv) Crucially, this entire collaborative workflow embeds robust security mechanisms - including granular permission control, dynamic security negotiation, and full-link auditing - to ensure all automated actions are controllable, trustworthy, and compliant to regulations.

Looking toward L5 "full autonomy," the core capability undergoes a fundamental evolution beyond the collaborative swarm intelligence of L4. L4's paradigm simulates human team division of labour, where a multi-agent swarm collaboratively solves ultra-large-scale, cross-domain problems that are intractable for any single agent. In contrast, L5 represents the highest stage of autonomous evolution, characterised by self-iteration, self-programming, the ability to create new agents and to generalise for new unforeseen situations, very high reliability, policy-compliant performance, and advanced capabilities for safety mechanisms and explainability. At this stage, the network transcends a pre-defined collaborative system and becomes an organic, self-evolving entity. Agents can autonomously discover novel fault patterns, adaptively generate and refine solutions through mechanisms like unsupervised learning within digital twins, and even dynamically program or spawn new specialised agents to address unprecedented challenges or architectural changes. This signifies the ultimate transition from intelligent collaboration to autonomous creation, where the network possesses genuine self-awareness, self-optimisation, and self-evolution capabilities. However, MNOs may still require a declarative approach to controlling, steering and trusting future L5 system behaviour in accordance with human needs.

3.2 SERVICE ASSURANCE

Business value:

Service assurance is a high-value scenario because it directly impacts customer experience, service reliability, and business outcomes. Operators increasingly need to warrant the performance of critical services in specific locations and time periods, such as mobile payments at stadiums, exhibitions, or major events. AI-driven service assurance can reduce service disruption, improve operational efficiency, and accelerate movement toward autonomous networks.

Current state and limitations:

Current service assurance processes are typically based on a combination of performance monitoring

systems, analytics platforms, alarm management systems, optimisation tools, and manual engineering expertise. While individual domains may support automation, end-to-end assurance remains largely manual. Several limitations remain: complex root-cause analysis across RAN and Core domains, siloed automation and limited service-level awareness, dependence on predefined rules and workflows, and predominantly reactive rather than proactive operation (justified by lack of standardised methods to measure proactive action). These limitations can increase operational costs and prolong service degradation.

Level-4 target state:

At Level-4 autonomy, service assurance becomes an autonomous, closed-loop process capable of continuously monitoring service quality, diagnosing degradation, coordinating remediation, and verifying recovery with minimal human intervention.

As in the fault management scenario, a typical service assurance workflow relies on hierarchical delegation from a supervisory agent to domain-specific agents.

In the example shown in Figure 2, the workflow is triggered by a service assurance intent (1). The cross-domain service assurance agent delegates diagnosis tasks only to the relevant RAN and/or Core network optimisation agent(s), as needed (2). Within each domain, the optimisation agent acts as a supervisor agent and delegates tasks to sub-agents (3). Sub-agents may query and share information directly between them and invoke digital twins and tools (4). The RAN and Core network optimisation agents return proposed root causes and handling solutions, and if the proposed actions lead to conflicts or trade-offs, the cross-domain service assurance agent drives conflict resolution and determines the final handling strategy (5). If the issue cannot be resolved remotely, field operations may be triggered, where field engineers are assisted by AI copilots (6).

Agent collaboration patterns:

The overall collaboration model follows principles already illustrated in the fault management scenario, including hierarchical task delegation, direct information exchange between agents, centralised conflict resolution, and the invocation of external tools and digital twins. The distinguishing characteristic of service assurance is its persistent focus on service-level outcomes rather than discrete fault-resolution events.

Key transformation enablers:

Key requirements for the transition from L3 to L4 service assurance are:

- Support for registration and discovery between agents. Selection methods may include, but are not limited to:
  a. Selecting an agent via a registration centre,
  b. Selecting based on the agent's own local knowledge/information, or
  c. Selecting according to user-provided input.
- Support for bidirectional proactive communication between agents, including task assignment, tracking, and result reporting; supporting information subscription between agents and active reporting of subscribed information.
- Support for authentication, authorisation, and trusted communication between agents.

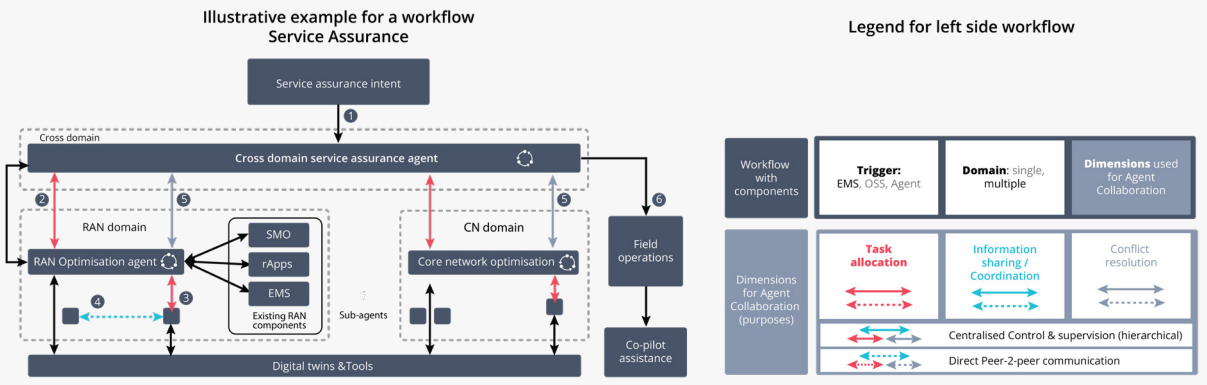

Note: The Cross-domain Service Assurance Agent invokes only the relevant domain agent(s). Agent selection may be based on, but is not limited to, registered scope and capabilities, available operational context, and user-provided information.

Figure 2: Illustrative example for a service assurance workflow

3.3 RAN OPTIMISATION

Business value:

RAN optimisation is a high value operational scenario because it directly drives service quality (throughput, latency, reliability, coverage), operational efficiency (reduced manual engineering effort), and cost outcomes (better utilisation of spectrum and infrastructure, and growing relevance of energy efficiency). Optimisation spans parameter tuning, traffic steering, interference mitigation, coverage/capacity balancing, and increasingly AI assisted control loops deployed across near real time and non real time control functions.

Operators view AI driven RAN optimisation as a key lever to progress toward zero touch / autonomous networks, while maintaining the required levels of trust, governance, and safety.

Current state and limitations:

Current RAN optimisation is commonly realised through a combination of SON style automation and analytics assisted engineering, with AI/ML increasingly introduced but constrained by practical operational requirements. Automation uses predefined rules, policies and workflows and leverages purpose-built tools. Typical implementations differentiate between near RT optimisation (10ms – 1s) and non RT optimisation (policy, model training, longer horizon optimisation), which already signals that timing, stability, and determinism are first order constraints in networks. Intent-based closed loop operation is already widely available leveraging e.g. RAN Service Management and Orchestration (SMO) and RIC environments with apps, rApps, domain controllers, AI/ML frameworks, data layers and telemetry.

Key limitations remain:

i) RAN optimisation has many interdependent parameters and objectives;

ii) trust and control requirements mean operators need explainability, governance and safe deployment;

iii) real time AI control must guarantee stability and predictable behaviour under strict timing constraints;

iv) RAN optimisation is today insufficiently integrated into E2E service management and network-wide optimisation of different KPIs incl. energy usage (e.g. rApps often optimise isolated metrics and lack multi-domain awareness);

v) today's systems are more rigid and reactive, and less proactive compared to agentic systems using LLMs/SLMs and predictive digital twins.

Level 4 target state:

At Level 4 autonomy, RAN optimisation becomes an ongoing, autonomous closed loop process (involving vertical loops across layers and horizontal loops within layers), continuously tuning the RAN for target objectives without waiting for scheduled engineering cycles, and with minimal human intervention beyond policy/intent definition and exception handling. The L4 target state is characterised by AI agents monitoring RAN telemetry, KPIs and inputs from other agents, and executing optimisation actions (e.g., load balancing, parameter tuning) autonomously with the help of typically SLMs or more traditional ML, while an upper coordination function ensures optimisation remains consistent with global objectives and avoids harmful cross effects.

Compared to L3 automation, L4 adds stronger intelligence and broader scope: autonomous optimisation decisions occur in real time or near real time, are guided by operator defined constraints, and can coordinate across multiple optimisation sub functions (e.g., interference, mobility, capacity, energy trade offs) as part of a policy bounded autonomous workflow. Due to the high level of deterministic, repeatable, policy-constrained and highly reliable RAN automation already in existence on levels 2-3, level 4 Agentic AI will likely be added to drive more intelligent self-sustaining autonomy, not to replace already sophisticated, specialised RAN domain controllers, RAN apps, services and workflow engines, which instead, where they exist, become available to AI agents as limited scope, already governed and validated 'precision tools' exposed to agents via MCP [32] or APIs. While AI agents will reason about complex decision challenges in non-real-time, rApps can run with bounded scope in near-real-

time. Thus, both can complement each other and enable high confidence in transitioning from L3 to L4 autonomy.

An illustrative example is depicted in Figure 3. In this example, the workflow is triggered by changes in telemetry data or an input from a service assurance agent (1). On the service operations layer, a cross-domain agent is delegating intent and allocating tasks to a RAN-specific optimisation agent also sharing policies (2). Within the RAN domain, an optimisation agent in the resource operations layer acts as a supervisor agent and delegates a task to a sub-agent residing in the infrastructure layer (3). Subagents may query and share information directly between them (4). All agents can invoke APIs, tools, knowledge graphs and digital twins, which may logically operate on different levels like service operations, resource operations and in the infrastructure layer. Where RAN optimisation actions run into conflicts with e.g. network-wide energy management, the cross-domain agent guides conflict resolution with the relevant energy management agent (5) and may arbitrate trade-offs between performance and energy consumption goals. Existing and possibly vendor-specific proven RAN automation and management components may continue to support deterministic and reliable actions through exposing their capabilities via APIs or agentic tool interfaces to any L4-capable RAN optimisation agent, as illustrated in the figure.

Agent collaboration patterns:

The collaboration model broadly follows the patterns already discussed in the fault management and service assurance scenarios. The main difference in the RAN domain is the coexistence of agentic decision-making with deterministic optimisation functions already operating in near-real-time and non-real-time control loops.

Key transformation enablers:

Achieving L4 RAN optimisation requires:

i) data availability and telemetry foundations (so optimisation is evidence based and timely);

ii) standardised contract-based interfaces that allow portable, vendor interoperable optimisation actions (also aligned with open RAN trends);

iii) appropriate compute placement for the required control loop timescales;

iv) in addition, operationalisation depends on governance, explainability, and safe AI deployment so the operator can trust autonomous optimisation ("no KPI degradation");

v) robust engineering to guarantee stability and determinism under near real time constraints. L4 systems often benefit from "agent tool interaction" where agents can validate candidate optimisation actions (e.g., via simulators or controlled environments) before applying them at scale across a wider RAN area.

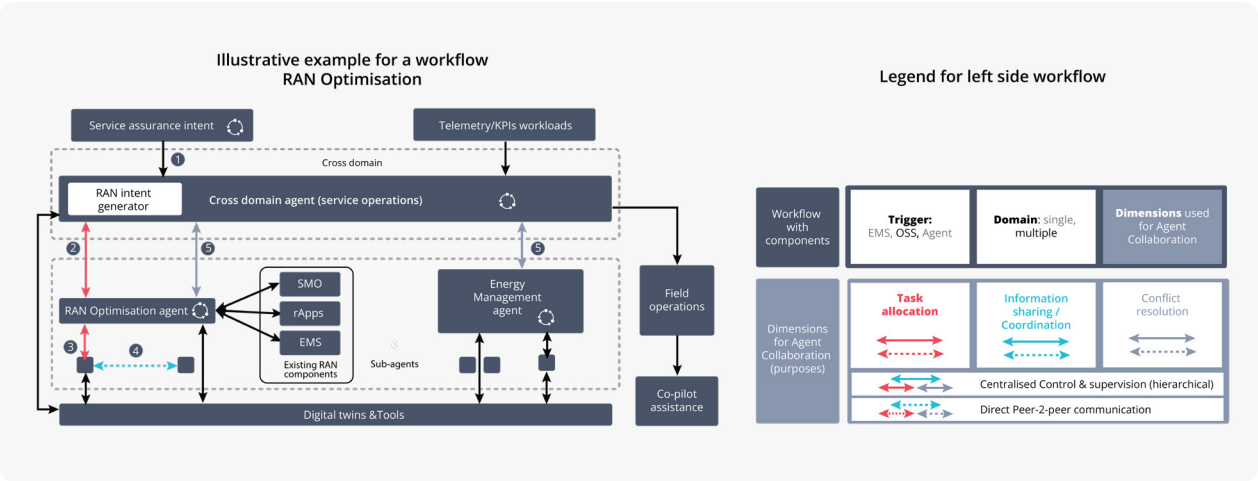

Figure 3: Illustrative example for a RAN optimisation workflow

04 REFERENCE FRAMEWORK FOR AGENTIC AI COLLABORATION

4.1 CORE COMPONENTS

The purpose of this section is to highlight components which, from a network operator perspective, are important for network automation with the help of Agentic AI and for moving from Level 2-3 autonomy to Level 4 autonomy. Such components eventually need to be in place, often supported by international standards, widely accepted open-source technologies and of course implementations. The document adopts a pragmatic approach, balancing high-level perspectives with sufficient detail to provide a clear and informative understanding of the key subcomponents and underlying concepts. Figure 1 presents the component categorisation framework that was selected.

While the columns are self-explanatory in Figure 4, the rows deserve some explanation.

The top row focuses on components which are important to enable human operation engineers and service provider personnel to effectively interact with an autonomously acting, automated network. Many of these required interfaces (mostly in the form of management dashboards) do not exist today but they will be required to let engineers communicate intents, intervene in exceptional circumstances (as human-on-the-loop), or approve decisions proposed by agents (as human-in-the-loop). Many components in the second row will also expose data to human personnel, where human oversight or intervention is required for aspects like explainability, observability, auditing, lifecycle management and accountability.

The second row (Agentic AI fabric) is a collection of aggregate components which are at the core of an autonomous, automated network, expected to continuously operate safely and reliably. This includes, of course, AI agents and the tools and AI models they have access to, AI Agent development, validation and verification systems (e.g. in form of digital twins), and several other components.

The third row is all about components related to a data fabric, which enables the autonomous operations. These components range from basic network instrumentation, telemetry systems, data collection and processing to knowledge representation suitable for Agentic AI systems (e.g. through knowledge graphs, ontologies, semantic data models).

The bottom row represents the resources and environments which from the perspective of AI agents are already there in the network. These existing network and IT resources, together with AI agent development platform, underpin the operations of autonomous agents in multiple ways. For instance, some agents may have access to small language models which can be trained using an operator's own AI platform. Domain-specific supervisory agents may also interact with network resources and non-agentic systems like domain controllers or even element management systems.

	Agentic Architecture	Agentic Validation	Agentic Action Observability	Agentic Lifecycle Management	Agentic Governance and Trust	Agentic Security
Human-Agent Interaction	Intent translation engine HITL, HOTL interfaces Agentic co-pilots	Intent validation	Agentic action monitoring dashboard Human action logging	Agent catalogue LCM Co-pilots LCM	Governance dashboard Tokenomics Guard Rails and Policies	Security dashboard
Agentic AI Fabric	Intent translation engine Agentic cross-domain orchestrator Agentic intra-domain orchestrator Intra-Domain agents Standard agentic protocols Agentic AI Capability Registration (Agent, Tool, AI Model, Skills, etc.) MCP clients, servers LLM/MCP gateways	Agentic validation system Agent test suites and sandboxes Digital Twins coordination agent Digital Twins repository	Agentic action logging Performance/ KPI monitoring Tool monitoring Resource usage monitoring Auditing of Agents logging	LCM for agents LCM for Agentic AI Capability (Agent, Tool, AI Model, Skills, etc.) Agent Mode Control DevOps/ MLOps/ LLMOps	Agent run time assurance Agentic behaviour analysis Agentic Audit System Policy repository Token metering and accounting tool Reputation	Agent interface security Agent authorisation Agentic ID management Agent software security Agent Access management Agent credential provisioning
Data Governance for Agentic AI and Resources	Data products, catalogues	Data sourcing and processing	Knowledge representation	Data mgmt and governance	Data strategy	
Resources and Environment	Network and IT resources	Training and Documentation	MNO AI platform	Infrastructure for Agentic AI	Agentic AI supply chain	

Figure 4: Overall framework of core components – Network Automation with Agentic AI

4.2 AGENT COLLABORATION PATTERNS

Mobile network operations are evolving from domain-specific automation towards end-to-end autonomous networks across differentiated services, network slicing, cloud-native functions, and exposed network capabilities. This evolution increases coordination complexity across RAN, core network, transport, cloud, edge, security, and service platforms. Without common interaction patterns, the introduction of multiple agents may create a new layer of fragmentation, including duplicated actions, unclear responsibility boundaries, and inconsistent levels of autonomous network capability.

From an operator perspective, agent interaction needs to address two questions: why agents interact, and how agents interact. From the purpose perspective, agent interaction can be broadly categorised into task allocation, information sharing and coordination, and conflict resolution. From the interaction style perspective, each category may be realised through centralised control and supervision, hierarchical coordination, or direct peer-to-peer communication.

Task allocation focuses on how operational responsibilities are divided. In selected scenarios such as fault management, service assurance, or RAN optimisation, an operational objective may need to be decomposed into data collection, diagnosis, impact assessment, solution generation, validation, and controlled execution. The main consideration is to clarify task scope, responsible party, and expected outcome, while avoiding duplicated or fragmented execution. Complex or cross-domain tasks are typically better supported through centralised control and supervision or hierarchical coordination, while local tasks may be handled through peer-to-peer communication for task requests and results.

Information sharing and coordination focuses on how agents form a consistent operational view. Cross-domain issues cannot be addressed reliably through local reasoning or raw data exchange alone. Agents need to share task context, affected objects or services, diagnostic evidence, analysis results, confidence levels, risk assessments, execution result, and capability catalogue for task-oriented discovery and matching. Semantic consistency is particularly important in cross-vendor and cross-domain environments. For large-volume or multimodal information, metadata and resource identifiers may be exchanged first, with detailed content retrieved on demand. During execution or after completion, feedback from a consumer agent to a producer agent is based on the monitored execution status or outcome of the requested task and supports the producer agent's subsequent task execution and coordination.

Conflict resolution focuses on operational risks caused by inconsistent objectives, conclusions, or actions. Conflicts may arise, for example, between performance optimisation and energy efficiency, between service continuity and security isolation, or between different root-cause assessments from different domains. Conflict resolution needs to be guided by operator-defined policies, including objective arbitration, negotiation and replanning, approval of high-risk actions, and escalation to operations engineers when necessary. Conflicts with limited impact may be handled through peer-to-peer negotiation, while cross-domain, high-risk, or policy-sensitive conflicts typically require coordination by an upstream agent or an agreed operational governance mechanism.

Overall, these three interaction patterns clarify how operational responsibilities are divided, how operational context is shared, and how inconsistent decisions are handled. They help support the evolution of network operations from isolated automation towards coordinated, end-to-end autonomous networks.

4.3 INTERFACE REQUIREMENTS

To make Agentic AI capabilities reusable across different operators, suppliers, and autonomous network scenarios, clear interface boundaries are needed among users, agents, tools, and network systems. The purpose of this section is not to define specific protocols, APIs, or implementation-layer specifications, but to identify the key functional boundaries required for interoperable and controllable agent collaboration. Interfaces can be categorised into four types: northbound interface, collaboration interface, southbound interface, and management interface.

The northbound interface represents the boundary between human or business intent and agent execution. It enables operations engineers, field engineers, and upper-layer systems to provide

operational objectives, constraints, and approval decisions to agents. The key consideration is to preserve operator control: agents need to understand and process incoming intents, which may be analysed and decomposed into sub-intents and recommended actions, while business priorities, risk tolerance, and approval rules should remain under operator governance.

The collaboration interface represents the boundary among agents. It enables agents from different network domains, suppliers, or operational roles to collaborate without relying on proprietary or scenario-specific integration. Its main role is to make collaboration relationships explicit: agents need to understand each other's roles, be able to trust each other, only interact within a permission scope, exchange necessary context, coordinate tasks, provide execution results, support feedback from consumer agent to producer agent, and escalate when issues cannot be resolved locally. This boundary is particularly important to avoid creating new silos in multi-agent deployment.

The southbound interface represents the boundary between agents and the systems or agentic tools they invoke, via agentic tool interfaces like MCP [32] or traditional protocols and APIs. These systems and tools may include OSS/EMS, network controllers, data platforms, digital twins, automation scripts, and AI services. The key consideration is controlled execution. In particular, tool invocation for actions that may affect service continuity, network stability, or security should be pre-validated (e.g. via digital twins), authorised , traceable, and reversible where necessary.

The management interface represents the boundary between agents and management functions that maintain agent lifecycle, operational status, and capability descriptions. It enables agents, agent platforms, or supervisory functions to register, update, and query agent profiles and capability catalogue, including agent roles, supported skills, applicable domains, invocation conditions, versions, availability, and policy constraints. The key consideration is governable capability management: capability information should be authenticated, consistent, up to date, and auditable, so that agents can be discovered, selected, and invoked for appropriate tasks under operator-defined governance.

Overall, the four interface types separate user intent, agent collaboration, tool invocation and management. This separation helps reduce integration complexity while maintaining operational controllability.

4.4 SECURITY AND TRUST MECHANISMS

Operators will adopt network automation with Level-4 autonomy only if they can trust the AI as much as a skilled human operator, and if they are confident that robust security mechanisms protect their networks. Security and trust mechanisms are treated in greater detail because MNOs consistently highlight them as critical adoption enablers.

- **Agent identity, role, permission control and task authorisation:**

  Every AI agent must have a provable identity, a tightly scoped role and act within the principle of least privilege. The identity should remain stable across the agent lifecycle so that verified operational history and reputation information can be consistently associated with the correct agent. Access should be granted just in time and should be short lived with e.g. built-in timer expiry for the access. No permissions should be lingering in the system after an activity is completed. An agent must explicitly break an intent down into specific technical actions (task identification) and prove it has the security rights to execute them (task authorisation), thereby preventing 'agent drift'. This protects the network from unauthorised or rogue actions while facilitating accountability.

- **Agent security interaction mechanism:**

  This mechanism enables agents to pre-negotiate and dynamically update security interaction rules required for collaborative task execution, such as security algorithms and corresponding parameters for source authenticity, content integrity, result verification, etc. In addition, the mechanism explicitly indicates whether dynamic updates of the above security context via re-negotiation during a task-execution are supported.

- **Policy governance and guardrails:**

  Pre-defined policies and guardrails constrain agent behaviour, preventing unsafe decisions even when agents operate without a human in the loop. A policy engine should sit in-between input and output guardrails. Before the policy engine even processes a request, an input guardrail checks it for safety, validity, and security. Even if a policy engine makes a logical decision, Agentic AI can sometimes generate hallucinated commands. The output guardrail is the final check before changes hit live infrastructure. Policy drift (when an AI agent or the network slowly shifts away from a desired state) should be identified e.g. via anomaly detection and remediated.

- **Secure-by-design agent architecture and infrastructure:**

  The AI agent platform itself must be secure. This means hardening the agent infrastructure (e.g., using secure coding, input validation, resource usage limits, robust authentication, encryption, and zero-trust principles for agent communications) to prevent exploitation by adversaries. Supply chain integrity of AI models and frameworks is also critical, given the threat of poisoned training data or malicious code injection. Threat detection via traditional and AI-driven security analytics and automated quarantine procedures for agents are further crucial security mechanisms. This must be complemented with data and model security (from training data, models, to knowledge graphs). Agent onboarding mechanisms must be robust, including sandbox testing for each agent before deploying it into a multi-agent system.

- **Reputation mechanism:**

  A dedicated reputation module establishes initial agent identities and baseline reputation, issuing verifiable reputation credentials that bind an agent to its historical behaviour. Reputation is assessed through a standardised scheme of digitally logged commitments and fulfilment feedback, yielding multi-dimensional, dynamic scores. The scoring framework supports time-decay, weighted aggregation of multi-role evaluations, and pluggable rules to adapt to different trust scenarios. Reputation scores directly govern service selection and access privileges, and tiered thresholds incentivise good behaviour while enabling progressive autonomy, thereby establishing a fair, practical ecosystem for trusted interactions among autonomous agents.

- **Pre-action verification and simulation:**

  Before agents execute changes on live networks, their proposed actions should be validated in risk-free environments using network Digital Twins for scenario planning and impact analysis.

- **Continuous monitoring and anomaly detection:**

  During operation, continuous oversight of agent behaviour, decisions and network conditions is essential. Telemetry and specialised anomaly detection systems need to verify that an agent's behaviour stays within expected bounds, detect behaviour anomalies in real-time and if required trigger automated containment and roll backs.

- **Dynamic trust and performance scoring with adjustment of agent responsibility:**

  Agents should be assigned a dynamic trust and/or performance score based on their performance and behaviour over time. If an agent consistently performs well, it can earn more autonomy and greater responsibility. This requires suitable KPIs for AI agent performance (e.g. decision accuracy, success rate of changes, false positive rates).

- **Gradual trust phasing mechanism:**

  A practical trust mechanism is to phase the introduction of autonomy. Network operators may initially run L4-capable systems in a supervised mode (effectively L3, with human oversight, or in a canary mode), gradually increasing autonomy as trust grows. This amounts to graduating agents to more responsibility over time. While this phased approach is certainly a process, it will require dedicated IT system support.

- **Continuous learning with guardrails:**

  Any learning must itself be monitored and controlled. While continuous learning makes

the system more adaptive, it poses challenges as AI agents change themselves. An agent that previously passed a certification may, after learning and self-adaptation, not pass it anymore.

- **Explainability, traceability and auditability:**

  To cultivate trust, agents must provide transparency into their decisions by producing rich logs and traces. This enables post-action examination and audits. Explainable AI techniques should be employed so that engineers understand why a particular action was taken.

- **Human override and fail-safe controls:**

  "Human-in/on-the-loop" capabilities (oversight consoles and dashboards, kill-switches, fallback modes) are required so that operations engineers can intercede if needed. When a kill-switch is activated, the system must not crash instantly. Instead, it must execute a controlled, orderly shutdown sequence that protects the live network and saves all operational data including logs. Kill-switches can be activated manually or automatically when red lines or thresholds get breached (e.g. KPIs, agent runaway behaviour, security breaches).

Below illustrates the use of above mechanisms following the phases of an AI agent's action lifecycle.

Pre-Action Phase: Secure Authorisation, Simulation and Policy Validation

In the pre-action phase, AI agents formulate or receive intents and plan network changes. Security and trust mechanisms here focus on ensuring that only valid, safe, and authorised actions can proceed beyond planning. Important in this phase are:

- Identity, role and permissions verification, with consideration of the agent's current trust status and reputation credentials (containing an aggregated reputation score and multi-dimensional scores)
- Network policy and service-level constraints compliance
- Digital twin simulation, dry-runs and sandbox testing
- Pre-deployment testing and agent certification especially for critical changes (an additional human approval or automated certification step might be required).

In-Action Phase: Runtime Monitoring, Enforcement and Anomaly Response

During the in-action phase, AI agents are actively executing changes or making on-the-fly decisions in the live network. The emphasis here is on real-time trust assurance. In this phase the following is important:

- Continuous behavioural monitoring of agents
- Policy enforcement points
- Dynamic trust scoring and adaptation of agents' responsibility accordingly
- Anomaly and rogue agent detection
- Automated containment and fall-back procedures

Post-Action Phase: Traceability, Audit and Continual Assurance

After an agent completes an action or a closed-loop cycle, post-action mechanisms provide transparency and accountability. Important in this phase are:

- Comprehensive logging and generation of traces
- Explainability and justification demonstration
- Post-action audits and reviews, with relevant outcomes used to update the agent's trust status and reputation credentials
- Continuous learning with guardrails

By rigorously accounting for above mechanisms, network operators can unlock efficiency and agility benefits while maintaining safety, reliability, and confidence.

While there is already significant industry consensus on mechanisms like fine-grained access controls, pre-deployment simulation and testing, monitoring and anomaly detection, auditability and compliance logging, there are some security and trust aspects still under debate or evolving:

This includes e.g. i) self-learning and adaptation allowing agents to self-modify their behaviour through machine learning in production as it complicates verification and control (see e.g. [52] on inter-agent generational knowledge transfer)

and ii) trust metrics and certification: The industry is still exploring how to measure and certify trustworthiness of autonomous networks. Concepts like "AI performance indicators" are still under debate. There's discussion about establishing certification regimes or "AI agent driver's licenses" to ensure solutions meet safety standards.

4.5 ALIGNMENT AND GAP ANALYSIS WITH EXISTING STANDARDS

Existing developments in alliances, SDOs and open-source organisations were analysed. Here a summary of the more detailed findings as presented in Appendix B. The questions were: Who does what, are efforts complementary and aligned? Most important, are there any areas that are not sufficiently addressed and that will concern network operators? These areas are opportunities for future concerted industry efforts.

Illustrative examples of relevant ecosystem activities considered during this study

There is a proliferation of work on agent-based AI for networks. Not all efforts narrowly target network automation, but many efforts are foundational to introduce Agentic AI. On the positive side, efforts in many SDOs and open-source communities are complementary as they address special topics related to Agentic AI for networks.

However, two issues are to be highlighted:

i) in some cases, work appears fragmenting;

ii) several areas important to operators are insufficiently addressed, representing critical gaps that should be addressed in the months to come.

A range of frameworks, guidelines, and specifications from various organisations that are relevant to network automation through Agentic AI were reviewed. The detailed findings of this review are provided in Appendix B. While these initiatives are contributing significantly to the advancement and maturity of Agentic AI for network automation, several important challenges remain. A summary of the perceived key gaps and unresolved issues is presented below.

Potential shortfalls, gaps and blind spots

1. **Fragmentation across organisations.**

   Many SDOs, alliances and open-source initiatives are community focussed and targeted, but sometimes also siloed. A unifying E2E perspective is missing. In Appendix B we provide examples related to 3GPP, TM Forum, ETSI, GSMA, Linux Foundation, AI-RAN Alliance and IETF. BBF could be added to this list, though it is more focussed on fixed access, thus relevant for converged service providers. A tangible example relates to agent-to-agent protocols.

2. **Limited focus on multi-vendor interoperability and agent coordination across domains.**

   We note the lack of a widely adopted, telecom-ready industry-wide agent communication protocol and, importantly, shared semantics. Today this makes true multi-vendor autonomy impossible. However, promising work is under way with the A2A protocol and extensions for telecom (A2A-T [8]), although having another variant of A2A [29] might be controversial.

3. **Lack of standards for agent knowledge- and context-sharing mechanisms.**

   The gap relates to mechanisms through which agents acquire knowledge and context and then share relevant context and situational awareness with other agents. Agents in different domains may work at cross purposes due to inconsistent situational awareness.

4. **Incomplete information modelling for networks.**

   We note a lack of unified information architecture to ensure agents across domains understand each other and interpret data and context consistently. We advocate for cross-SDO alignment and shared ontologies.

5. **Lack of end-to-end security, identity and trust for autonomous functions.**

   Present frameworks do not define how to handle unique agentic AI security risks, how to manage agent identity and safe operation across domains. We advocate for a Zero-Trust Agent Ecosystem supported by unifying specifications and standards.

6. **Limited focus on the human-machine interface and operational integration.**

   Frameworks do not sufficiently address how human network engineers receive escalations, validate, override, or tune agent decisions at runtime. This risks inconsistent safety practices.

7. **Lack of operator-friendly governance and lifecycle tools.**

   Detailed frameworks for governing autonomous AI agents across their lifecycle are missing. They need to address aspects like agent onboarding, certification and testing, common performance indicators, agent maturity graduation schemes, continuous agent assurance and performance management.

8. **Under-addressed economic and organisational readiness.**

   There is lack of guidance on required operator organisational changes and cost management support frameworks allowing operators to control LLM token consumption and to strike a balance between operational safety, performance and cost associated with using agentic AI.

9. **Agentic AI architecture when network scales.**

   There is limited guidance on how Agentic AI solutions should scale across large telecom environments, including hierarchical agent architectures, workload partitioning, distributed reasoning, and the use of graph-based AI techniques such as Graph Neural Networks (GNNs) to efficiently model and reason over complex network relationships.

10. **Lack of guidance on multi-agent system sising based on use case complexity.**

    There is no standardised methodology for determining the appropriate number of AI agents based on network size, operational complexity, service criticality, or latency requirements. This creates uncertainty in balancing system performance, reasoning quality, operational overhead, and deployment cost.

11. **Limited support for low-code and no-code agent development.**

    There is no standardised development methodology or modelling language enabling operators and domain experts to design, configure, test, and maintain Agentic AI workflows without extensive software development expertise, thereby slowing adoption and reducing portability across platforms.

12. **Lack of abstraction layer for closed-loop automation.**

    Existing closed-loop automation frameworks do not adequately define a vendor-neutral abstraction layer capable of translating heterogeneous vendor-specific alarms, KPIs, configuration parameters and management interfaces into a unified semantic representation consumable by AI agents operating in multi-vendor networks, that could make agentic systems more modular and pluggable solutions.

Appendix B provides more detail regarding both what's going on in different organisations and where specific, perceived gaps exist. Many of these gaps are explicitly covered in our reference framework.

05 CONSIDERATIONS FOR THE ECOSYSTEM

5.1 PURPOSE AND POSITIONING OF CONSIDERATIONS

The ecosystem relevant to this report includes MNO´s, vendors, solution providers, AI and cloud service providers, standards-defining organisations (SDOs), alliances and open-source communities.

The considerations in this chapter are articulated from the perspective of mobile network operators. They are not intended to prescribe how other organisations should conduct their work. Rather, they highlight operator requirements, cross-ecosystem dependencies and gaps that need to be addressed if Agentic AI is to become a safe, secure, interoperable and scalable enabler of higher levels of network automation.

The analysis conducted for this report indicates that substantial work is already under way across SDOs and open-source communities (e.g. TM Forum, 3GPP, ETSI, W3C, IETF, AI-RAN Alliance, BBF and Linux Foundation communities). This is encouraging. At the same time, the distribution of work across many organisations creates risks of fragmentation, inconsistent assumptions, duplicated effort and gaps in operator-critical areas. The considerations below therefore focus on areas where early alignment would increase confidence, reduce adoption risk and accelerate progress towards network automation in Level 4 autonomous networks.

In addition, closing gaps through standards or open-source helps the industry without limiting vendor differentiation. Common interfaces, trust mechanisms and lifecycle models should create a larger market for deployable agentic automation. Vendors can still differentiate in models, AI agents, optimisation algorithms, domain expertise, performance, usability and operational outcomes.

The conceptual framework discussed in Chapter 4 serves here as a high-level checklist, along the lines “sooner or later, many or all of these core framework components will become critical, thus, they need to be on our radar screen”. NGMN encourages the ecosystem to take into consideration the requirements of the proposed framework and reflect on whether those are sufficiently addressed in their respective areas of work currently or need further addressing. NGMN has identified the following generic considerations for SDOs, the open-source and vendor community which are relevant to support MNOs in developing an agentic framework. More details are provided in Appendix B.

5.2 CONSIDERATIONS FOR OPEN-SROUCE COMMUNITIES, STANDARDS AND INDUSTRY ORGANISATIONS

The following are opportunities that could help the mobile networks industry.

1. Strengthening cross-SDO and open-source coordination on Agentic AI for network automation. Relevant organisations to establish lightweight but recurring cross-community coordination mechanisms, such as shared gap registers, and cross-referenced work-item maps.

2. Actively resolving deployment-critical overlaps. This can be done through joint special sessions or workshops between organisations where overlapping proposals are compared against operator requirements such as multi-vendor interoperability, security, latency, deterministic behaviour and others.

3. Mapping blueprints and reference architectures to realisation patterns. This can be done through publication of implementation-neutral realisation patterns and/or any emerging solution proposals, whether proprietary or standardised.

4. Developing telecom-grade agent communication profiles with shared domain semantics where generic protocols are deemed insufficient.

5. Standardising principles for context sharing, context scoping and situational awareness. This to include context relevance, dynamic scoping to meet cost control and enforced token budgets, summarisation, freshness, provenance, confidence, data minimisation, privacy, domain boundaries, memory retention, and expiry.

6. Aligning information models, ontologies and semantic mappings across network domains. The industry should define common core concepts, domain-specific extensions, and mapping mechanisms between TM Forum, 3GPP, ETSI, O-RAN, IETF, W3C, BBF and open-source models.

7. Defining a telecom-grade Zero-Trust Agent Ecosystem. Such an ecosystem should address aspects like agent identity, authentication, authorisation, secure communication, tool access, workflow verification, software provenance, runtime attestation, policy enforcement, least-privilege access, credential lifecycle management, audit logging, behavioural monitoring, revocation, quarantine, incident response and others.

8. Establishing interoperability testing, certification and plugfest activities.

9. Treating safe execution, fallback and rollback as first-class requirements.

10. Defining requirements for simulation and digital-twin environments.

11. Standardising observability and auditability for autonomous agents.

12. Providing migration guidance from existing automation to agentic AI-enabled autonomy.

5.3 CONSIDERATIONS FOR MNOs

1. Defining operator requirements for human-agent interaction and operational integration.

2. Establishing common operator requirements for agent lifecycle governance and assurance.

3. Developing operator guidance for organisational readiness, operational efficiency and resource-aware Agentic AI.

4. Collaborating on test scenarios and acceptance criteria for high-value autonomy use cases.

5. Sharing non-sensitive operational requirements and lessons learned with SDOs and open-source communities.

6. Addressing data readiness, telemetry quality and data governance.

5.4 CONSIDERATIONS FOR VENDORS, SOLUTION PROVIDERS AND HYPERSCALERS

1. Implementing open, profile-based and interoperable agent interfaces and ideally common data models.

2. Making controllers and management platforms safely consumable by (higher-level) agents.

3. Providing verifiable evidence of agent behaviour, safety and interoperability.

4. Supporting operator-controlled governance of AI agents with respect to behaviour, safety and interoperability.

5. Supporting hybrid, multi-cloud and sovereignty-conscious deployment models.

6. Contributing to shared validation environments, testbeds and plugfests.

As mentioned above, all considerations are elaborated in more detail in Appendix B.

06 CONCLUSION

The telecommunications industry is entering a new phase of network automation in which AI is no longer limited to analytics, prediction, and closed-loop optimisation. The emergence of Agentic AI introduces the possibility of autonomous systems that can reason, plan, collaborate, and execute actions within and across operational domains while continuously adapting to changing network conditions and business objectives.

This publication has assessed the motivation, opportunities, requirements, and ecosystem developments related to applying Agentic AI to network automation. From an MNO perspective, it contributes three key outcomes:

i. a pragmatic reference framework identifying the core components required for Agentic AI-driven network automation (Chapter 4),

ii. a brief structured assessment of ecosystem maturity and potential or perceived gaps across standards and open-source activities (Chapter 4), and

iii. a set of targeted considerations to accelerate safe, interoperable, and scalable deployment (Chapter 5).

The analysis indicates that Agentic AI will likely become a key enabler for achieving higher levels of automation and network autonomy, supporting the industry's ambition to progress towards Level 4 and ultimately Level 5 autonomous networks.

At the same time, the findings demonstrate that realisation of this vision requires significantly more than the capabilities currently emphasised in the wider AI ecosystem, such as LLMs, agent frameworks and tool integration. It will require robust and standardised mechanisms for agent registration, discovery, communication, task delegation, negotiation, conflict resolution, policy enforcement and shared situational awareness. Trustworthiness, security, explainability, observability, interoperability, governance and human oversight must be treated as foundational design principles as outlined in Section 4.4.

The review of SDOs, industry initiatives and open-source activities demonstrates encouraging momentum across the ecosystem. However, the current landscape remains complex and unevenly mature. In Appendix B, this report identifies several potential or perceived gaps. Addressing those could substantially help the mobile networks industry.

These gaps highlight that achieving autonomous networks is not primarily a technology integration challenge, but an ecosystem-level transformation requiring a pragmatic phased adoption (i.e. to protect ROI) and a coordinated progress across architecture, standards, security, operations and governance.

Furthermore, mobile networks present fundamentally different requirements in several areas compared with enterprise IT environments. They demand high reliability, deterministic behaviour, real-time responsiveness in critical domains, regulatory compliance and strict safety guarantees. As a result, general-purpose Agentic AI technologies must be adapted, extended and hardened where necessary to meet telecom-grade operational requirements.

Over the next few years, the industry is expected to transition from isolated proof-of-concepts towards operational deployments of Agentic AI in selected network domains, focusing initially on high-value scenarios where tangible business and operational benefits can be demonstrated. Early deployments will likely prioritise bounded and well-controlled use cases such as RAN optimisation, energy efficiency, fault management, incident resolution, service assurance, engineering support and cross-domain troubleshooting, where sufficient human oversight can be maintained (HOTL).

This transition is expected to follow a phased approach, progressing from supervised and recommendation-based systems towards bounded

autonomous operation and, over time, towards collaborative multi-agent environments. As confidence grows and trust mechanisms mature, agentic systems are expected to evolve from isolated task-specific agents into coordinated multi-agent ecosystems capable of managing complex workflows across network, service, cloud and business domains.

Apart from the reference framework, key takeaways of this report are the following identified opportunities for more industry work or intensified communication of already achieved results in those areas:

- Convergence across standards bodies and open-source communities to reduce risk of fragmentation and align on interoperable architectures, interfaces and responsibilities
- Development of a telecom-grade Zero-Trust Agent Ecosystem covering, for example, agent identity, authentication, authorisation, policy enforcement, complemented by auditability and explainability
- Standardisation of agent communication protocols, shared semantics and efficient context-sharing mechanisms across domains
- Alignment of information models, ontologies and data architectures to enable consistent cross-domain understanding and reasoning
- Establishment of lifecycle governance, certification, testing, evaluation and lifecycle assurance frameworks, including maturity models for AI agents
- Expansion of digital twins and simulation environments as mandatory validation layers for agent behaviour and decision-making
- Definition of operator-centric operational models, including human–agent interaction, organisational readiness, skills transformation and resource-aware operation of AI systems.

In parallel, MNOs, vendors, hyperscalers, standards bodies and open-source communities will need to collaborate closely to ensure that solutions remain interoperable, secure and aligned with operator requirements. Cross-industry coordination, shared testing environments, and evidence-based validation of agent behaviour will be essential to build confidence in large-scale deployments.

In conclusion, Agentic AI should therefore be regarded as a foundational capability for future telecommunications operations and network automation with several gaps still to close, rather than a standalone technology trend. The considerations in this report shall contribute to delivering a coherent, interoperable, trusted and sustainable ecosystem capable of delivering Level 4+ autonomous networks at scale.

APPENDIX A: GLOSSARY

Term	Definition
AI Agent	An autonomous or semi-autonomous software entity capable of perceiving inputs (e.g., telemetry, events, or intents), reasoning, making decisions, and performing actions to achieve defined objectives. The 'brain' of an agent can be a Large Language Model or a Small Language Model.
Agent Collaboration Pattern	A defined mode of interaction among agents, including e.g. task allocation, information sharing, and conflict resolution mechanisms, which determines how agents cooperate to complete tasks.
Agent Lifecycle Management	The processes governing the creation, deployment, configuration, monitoring, update, and retirement of agents within the system.
Agentic AI	An AI paradigm in which AI systems act as agents that can autonomously plan, reason, collaborate, and execute tasks, often in coordination with other agents and guided by human-defined intents, policies and guardrails.
Agentic AI Capability	A system-level operational capability created when multiple components like AI agents, LLMs, tools, data, management services, interfaces and governance mechanisms work together. Examples: ability to interpret operational objectives, reason over network context, plan and decompose tasks etc.
Agent-to-Agent (A2A) Interface	A specific type of collaboration interface enabling standardised communication and interaction between agents.
Agent-to-Tool Interface	An interface enabling agents to invoke external tools, APIs, or systems to perform actions or retrieve data.
Agent Reputation	A multi-dimensional, dynamic score bound to an agent's identity, issued as verifiable credentials by a trust registry based on historical commitments and fulfilment.
Autonomous Network (AN)	A network capable of self-configuration, self-optimisation, self-healing, and self-protection with minimal human intervention, guided by intent-based management principles.

Digital Twin	A virtual representation of a network, system, or component used for simulation, validation, and optimisation of actions before real-world execution.
EMS	Element Management Systems, is a software platform used in mobile networks to manage and configure specific network elements such as baseband units.
Feedback	A collaboration mechanism in which a consumer agent, after requesting a task from a producer agent, monitors the task's execution status or outcome and provides task-related feedback to the producer agent during execution or after completion.
HITL	Human in the loop. Refers to humans having a role to play inside a closed-loop automation system (e.g. for approval purposes).
HOTL	Human on the loop. Refers to humans only playing a role in exceptional circumstances in the context of closed-loop automation with AI agents.
Intent-Based Management	An operational paradigm where users or systems specify desired outcomes ("intents"), and the system automatically determines and executes the required actions to achieve them.
IRTF	Internet Research Task Force. An open international research organisation focussing on the evolution of the Internet and network technologies with concepts not yet mature for standardisation in the IETF.
IRTF-NMRG	Internet Research Task Force – Network Management Research Group. A research group within the IRTF investigating emerging technologies for network management, including topics such as self-driving networks, intent-based networking, and AI-driven management. Focuses on pre-standardisation research.
Knowledge Graph	A structured representation of knowledge (entities and relationships) used by AI agents to support reasoning, inference, and decision-making.
LCM	Lifecycle management
Level 3 Automation (L3)	A stage of network automation characterised by rule-based and partially AI-assisted automation, requiring significant human oversight for cross-domain coordination and decision-making. [7]

Level 4 Autonomy (L4)	A stage of autonomous networks where most operational processes are executed autonomously by AI agents within policy constraints, with humans involved primarily in exception handling and governance. [7]
Level 5 Autonomy (L5)	The highest stage of network autonomy, characterised by self-learning, self-evolution, autonomous creation of new agents or logic, and the ability to handle unforeseen scenarios with no or minimal human involvement. [7]
LF-AAIF	Linux Foundation – Agentic AI Foundation. An open, neutral foundation under the Linux Foundation that supports development and governance of open-source technologies, standards, and ecosystems for agentic AI.
LF-AGNTCY	Linux Foundation AGNTCY Project. An open-source infrastructure project under the Linux Foundation that provides foundational capabilities for multi-agent systems, forming a core layer for the "Internet of Agents."
LLM	Large Language Model
MCP	Model Context Protocol (for the interface between agent and tool)
MNO	Mobile Network Operator
Multi-Agent System (MAS)	A system composed of multiple interacting agents that collaborate, coordinate, or compete to achieve individual or shared goals in a distributed environment. MAS have already been in existence before the arrival of Agentic AI and Generative AI.
OpenID	OpenID Foundation. Creates identity standards that are secure, interoperable, and privacy-preserving.
OSS	Operations Support Systems
OWASP	OWASP® Foundation works to improve the security of software through its community-led open-source software projects.
Policy Constraints	Rules and boundaries defined by operators that guide and limit agent decision-making to ensure safe and compliant operations.
SLM	Small Language Model
TM Forum	Telemanagement Forum
Tokenomics	The discipline of managing, predicting, and optimising token consumption
W3C	World-Wide Web Consortium

APPENDIX B: EXISTING STANDARDS, OPEN SROUCE PROJECTS AND INDUSTRY INITIATIVES

In this Appendix we look at the status of the industry, areas of alignment across the industry and highlight potential gaps, blind spots and issues which, if confirmed, offer an opportunity for collective industry action.

Disclaimer: This appendix reflects the authors' understanding of selected standards, industry initiatives, and open-source activities related to network automation, autonomy and Agentic AI at the time of writing. Given the rapid pace of developments and the breadth of ongoing work, the analysis is neither exhaustive nor definitive. Observations on potential gaps, challenges, or opportunities are provided from a network operator perspective to encourage constructive industry discussion and collaboration, and should not be interpreted as criticism of any organisation, initiative, or specification. Some issues identified may already be resolved, under active consideration or resolution, and omission of specific activities does not imply lack of relevance.

Illustrative examples of relevant ecosystem activities considered during this study

Efforts toward **agent-based AI for network automation** have proliferated across standard-defining organisations (SDOs) and open communities, posing some issues for network operators.

TM Forum has emerged as a frontrunner, delivering a 6-level taxonomy for network autonomy (Levels 0–5) widely adopted across the world [7]. Some national standards associations like CCSA have already formalised it further into national industrial standards. Moreover, TM Forum produced e.g. a comprehensive Autonomous Networks framework (IG1218F) [4], an Autonomous Networks Level 4 target architecture (IG1251C) [5] and a guide on Multi-Agent Architectures (IG1251E) [6]. Dedicated projects like AI Native Blueprint [15] and Autonomous Networks [16] are under way. TM Forum has delivered detailed guides on multi-tier closed-loop designs (aligning with 3GPP and ETSI frameworks), agent-based architectures, and proactive AI-enabled operations. They also created maturity and evaluation frameworks (like IG1392 [10], IG1252 [7]) to benchmark progress.

3GPP has also advanced standards relevant to agentic network automation. 3GPP SA5 (Management, Orchestration and Charging)) integrated the TM Forum's AN definition into mobile network standards. Focus was placed on standardising the technical building blocks required to achieve AN Levels, including Intent-Driven Management and Closed-Loop Automation. Efforts on autonomous networks incl. AI agent interaction (e.g. 3GPP TS 28.100 [19]) and intent-driven management (e.g. 3GPP TR 28.312 [20]) lay groundwork for policy/intent-based agent orchestration. While largely confined to the management-plane, 3GPP's contributions provide a common vocabulary for autonomy levels and network intents that can unify industry understanding. More exploration into aspects of network autonomy and use of agentic AI including for network management is happening as part of the 6G study item (e.g. 3GPP SA1 with TR 22.870 [21], 3GPP SA5 with TR 32.801-1 [22], 3GPP SA2's AI4Net [KI#18] - key issue#18 on how AI is used inside the network to design, operate, and optimise it).

ETSI published foundational white papers and specifications mapping out how AI, machine learning, and cognitive technologies power the transition up through AN level. **ETSI ZSM** (Zero-touch network and Service Management [23]) defines an end-to-end, service-oriented autonomous management

architecture with hierarchical closed loops, which dovetails with TM Forum's approach. **ETSI ENI** (Experiential Networked Intelligence [24]) focuses on how AI/ML can bring policy-driven cognition to network management, e.g. contextualising decisions similarly to the role of "cognitive management" in 3GPP SA5's work. This addresses real-time orchestration across vendor domains, enabling evolution toward multi-agent approaches. **ETSI OCG-AN** (Operational Co-ordination Group on Autonomous Networks, [25]) is a cross-ETSI coordination body that aligns and orchestrates all ETSI activities related to autonomous networks (AN).

Various industry alliances and open-source communities are contributing key building blocks. We believe SDOs should be fully aware of those activities and look for synergies.

The **O-RAN Alliance** has delivered open interfaces and a RAN Intelligent Controller (RIC) architecture enabling plug-in "xApps/rApps" for automated radio optimisation [26], [27]. This effectively introduces specialised RAN agents into live networks today, demonstrating practical agent-based automation. This offers a foundation in the RAN domain to build on by higher-layer agents. The **AI-RAN Alliance** has a dedicated "Agentic AI Task Group" which aims to address autonomous network decision-making, self-optimising network operations, AI agent orchestration frameworks and intent-based network management [28].

The **Linux Foundation** hosts relevant working groups and open projects: for example, the **Agentic AI Foundation (AAIF)** backed e.g. by Google, Microsoft, AWS and OpenAI establishes open standards and interoperability protocols for multi-agent environments. The Linux Foundation is the home of the foundational **Agent-to-Agent Protocol (A2A)** as contributed by Google and Agent Communication Protocol (ACP) from IBM [29]. Within AAIF, the important **Model Context Protocol (MCP) Working Group** standardises how AI agents dynamically discover and consume automation tools, telecom telemetry and APIs without continuous code refactoring [30]-[32]. The LF project **AGNTCY** has work in progress on inter-agent discovery, secure messaging, identity, and tracing [34].

The **CAMARA** project defines how to expose 3GPP network capabilities as APIs that AI agents can discover and manipulate [35]. **LF project Nephio** uses Kubernetes-based intent-driven GitOps to act as a crucial safety execution layer for AI agents, so they do not cause network outages [36]. **LF project Sylva** provides the standardised , cloud-native telecom infrastructure needed to run resource-intensive AI workloads, specifically localised small language model reasoning engines [37]. **LF Edge** standardises the bare-metal and containerised environments hosting distributed edge AI agents [38].

Other forums are acknowledging agentic AI's importance. **GSMA's 6G Taskforce** published a white paper on "AI Agents & Agentic Protocols" in early 2026, highlighting major gaps and enumerating requirements for telecom-grade multi-agent systems (like goal-driven agent protocols, specialised semantics, and trust tools) [39]. **ITU's Focus Group on Autonomous Networks** (within ITU-T SG13, [40]) has started aggregating these industry insights to advise on pre-6G frameworks; though still nascent, it drives alignment among regional SDOs.

The **IRTF's NMRG (Internet Research Task Force - Network Management Research Group)** has effectively become the primary incubator for routing-plane and management-plane Agentic AI standards for the internet. In early 2026, several Internet-Drafts were introduced to specifically formalise how AI agents communicate leveraging A2A (e.g. a controller agent to a device agent sitting on a network element), how protocol interoperability and tool bridging can be achieved for agents interacting with legacy southbound network element interfaces [53]-[58].

In **IETF**, work is emerging on a trust framework, as autonomous agents cannot be allowed to operate anonymously. IETF is exploring extensions to cryptographic identity protocols (like OAuth, SPIFFE/SPIRE, or decentralised identifiers, see W3C) specifically for AI entities. IETF is also addressing how agent-to-tool protocols (like MCP) can be reliably layered over existing internet substrates (like TLS/WebSockets/HTTP3) in carrier-grade infrastructure. In June 2026, IETF approved a BoF request titled 'Agent Communication Protocols', which may result in a new working group with potential focus on agent-to-agent protocol, session protocols for AI agents and human-in-the-loop interactions [41]-[51].

While the **W3C** does not standardise telecom network architecture per se, its foundational standards are

critical for driving Autonomous Networks Level 4 (e.g. [61]-[66]). W3C's Semantic Web (SOSA/SSN) and Web of Things (WoT) frameworks provide unified ontologies needed for cross-domain intent processing, while its Verifiable Credentials (VC) [68] and Decentralised Identifiers (DID) [67] standards provide an automated, cryptographic trust framework intended for autonomous AI agents to negotiate and execute closed-loop actions safely.

While the **Broadband Forum (BBF)** is not directly relevant for mobile networks, parts of the mobile networks industry may benefit from a degree of alignment between mobile and fixed networks (e.g. converged service providers and vendors supporting both types of access networks). BBF has outlined a blueprinting roadmap for Agentic AI and AI Agents in the fixed access domain [59]. BBF finalised Phase 1 of this initiative by publishing a formal vision marketing report outlining the blueprint of an AI-powered internet [60]. This document established how AI agents will deliver intelligent user support and power self-optimising, self-healing networks. The working texts WT-536 and WT-537 represent ongoing Phase 2 work to standardise the integration of artificial intelligence (AI) into broadband networks. WT-537 specifically addresses the requirements, architecture, and behaviour of AI agents within the telecom infrastructure and appears therefore most relevant for the context of this report.

Because of this plethora of efforts, there is a degree of risk of fragmentation and partial duplication. We advocate for 'cross-silo awareness', open-mindedness, and the need to converge on unifying standards and open-source frameworks for the sake of scale and cost benefits.

Potential shortfalls, gaps and blind spots

We highlight potential gaps, issues and blind spots which, if confirmed, could be addressed by relevant organisations to meet the needs of network operators.

A) Fragmentation across organisations.

Despite encouraging progress, initiatives often remain siloed within each SDO or consortium, lacking a unifying end-to-end view from a service provider perspective. This is possibly a result of community dynamics and the sheer complexity of today's networks. This makes the compilation of a coherent end-to-end framework from a mobile network operators' perspective more difficult. Examples: 3GPP's autonomy levels focus on network management architecture, but they appear to stop at domain boundaries and seem to so far assume human oversight in cross-domain scenarios. TM Forum's comprehensive blueprint is not yet fully mapped to each domain's specifics and for now this is left to vendors for actual solutions. ETSI ZSM focuses on closed-loop patterns but, maybe intentionally, doesn't prescribe yet how intelligent agents should be structured or interact beyond functional blocks, possibly leaving an implementation gap. GSMA's cross-industry view on agentic AI is very valuable but so far stops short of creating actionable specs in this area. IETF's work on agentic AI is largely only emerging now but needs to be aware of important work in TM Forum, 3GPP, Linux Foundation and W3C. Finally, some friction may exist for work on agent-to-agent protocols (e.g. A2A [29] in Linux Foundation, A2A-T [8] in TM Forum, corresponding work in LF AGNTCY). While they approach the problem from different industries, they are all potentially independently building schemas, specifications, and protocols to handle agent capabilities, agent skills discovery, and intent translation. Above is not criticism, instead reflects how the authors perceive the overall landscape. We promote more harmonisation and collaboration in this area. Using foreseen extension mechanisms (as A2A-T does with A2A) helps in this respect.

B) Limited focus on multi-vendor interoperability and agent coordination across domains.

Current initiatives primarily address automation **within** a domain or system, but operators need coordinated autonomy across RAN, transport, core, cloud infrastructure and services. The lack of a widely adopted, telecom-ready **industry-wide agent communication protocol** is a major gap. The few emerging proposals (A2A as in Linux Foundation, TM Forum's version A2A-T, and others) are not yet unified. Without interoperable agent communication and **shared semantics**, true multi-vendor autonomy may well be impossible. We advocate for extensible agent interfaces (to agents and tools) that allow plug-and-play of AI "agent apps" across networks, and robust inter-agent protocols that support structured intent-sharing, negotiation and task delegation together with predictability and semantic grounding, as required for telecom operations. A single widely adopted

approach fit for telecom purposes is desirable to produce scale and cost benefits.

C) Lack of standards for agent knowledge- and context-sharing mechanisms.

Little standardisation appears to exist for **how** agents in different systems **share situational context** or "world models" (in contrast to what to share for agent understanding, see D). For example, 3GPP's management standards specify telemetry collection functions like NWDAF, but do not define **how** an AI agent would consume that data, integrate external knowledge repositories or consume insights from digital twins to enhance its own knowledge and context and share with another agent where appropriate. TM Forum's AI Data Architecture guide (e.g., IG1356 [3]) highlights the need for modern data fabrics and knowledge graphs but remain advisory frameworks, not yet widely adopted specifications. This potential gap risks agents in different domains working at cross purposes due to inconsistent situational awareness.

D) Incomplete information modelling for networks.

It appears there is **no sufficiently unified information architecture** to ensure agents across OSS, RAN, Core, etc., interpret data and context consistently. Today's standards provide conflicting or incomplete data models – e.g., 3GPP's network management models vs. TM Forum's Information Framework/SID can diverge on naming and scope; vendors have their own closed-source vendor schemas. Without harmonisation, common ontologies enabling knowledge graphs/reasoning graphs and an AI-friendly data fabric, an agent in one domain cannot fully understand metrics or events from another, limiting cross-domain intelligence. Distributed knowledge should be consistent as it is used for reliably 'grounding' telco AI. TM Forum's "Data Architecture for AI-Enabled Operations" [3] guide notes that telecom data is heavily siloed and heterogeneous, urging adoption of a common ontology. We advocate for cross-SDO alignment and shared ontologies.

E) Lack of end-to-end security, identity and trust for autonomous functions.

Beyond general IT security measures, present frameworks seem not to define in great detail how to handle unique agentic AI security risks like prompt injections, malicious or errant agents, or novel failure modes in multi-agent systems. Further, the authors understand there is no standardised framework yet to manage agent identity, authorisation, and safe operation across domains. Mobile networks cannot cede control to AI without absolute clarity what each agent is allowed to do, and how its decisions are monitored. Network operators will benefit from an architecture for agent safeguards with runtime oversight, kill-switches, rigorous testing regimes and agent maturation/graduation lifecycle schemes (from low-trust/canary deployments to high-trust operational modes) – none of which seem to be fully covered by current standards. TM Forum's guide on Agentic Security (GB1087) [2] acknowledges this situation and TM Forum IG1463 Security Considerations for Agentic AI in Autonomous Networks [9] lists several Agentic AI security risks. We advocate for the creation of a **Zero-Trust Agent Ecosystem**, including standards to define how agents prove identity, get "licensed" for network actions, demonstrate their growing maturity, log decisions for audit, and gracefully hand control back to humans when necessary.

F) Limited focus on the human-machine interface and operational integration.

Standards typically assume that if autonomy is high, human involvement is minimal, but network operators need clarity on how humans remain in or on the loop for oversight and exceptions. Efforts by TM Forum and 3GPP to define human "intent" interfaces are essential, but beyond that, frameworks do not deeply address how human network engineers receive escalations from agents, validate, override, or tune agent decisions at runtime. Without clear guidance, operators might develop ad-hoc solutions, risking inconsistent safety practices. TM Forum has begun specifying new operational roles and processes for autonomous networks, but other organisations have yet to incorporate how human oversight and safety guardrails will be managed across agentic AI-driven systems.

G) Lack of operator-friendly governance and lifecycle tools.

The industry still lacks detailed frameworks for governing autonomous AI agents across their lifecycle and for meeting regulatory requirements. Network operators need confidence that they can rigorously

manage AI/ML models and AI agents – from model validation, agent certification and testing through to continuous agent updates, moving agents after a qualification period to a higher, more responsible mode of operation (via a graduation scheme), continuous agent assurance and performance management. TM Forum's AI Native Blueprint project [15] touches on concepts like agent onboarding, sandbox testing, and drift handling, but other SDOs have been mostly silent. In addition, standardisation blind spots seem to exist related to auditability and explainability of agent decisions. To close this perceived gap, we advocate for the development of cross-industry best practices and/or dedicated workstreams.

H) Under-addressed economic and organisational readiness.

While technology groups focus on architecture, few have formal guidance on the required telco organisational changes (except TM Forum's emerging work on AN governance capabilities) and cost management support frameworks (e.g. metrics and measurement functions for agent token consumption, cost prediction for LLM usage and related real-time cost controls, which overall might be termed 'tokenomics' [73]-[77]). The risk is that even if technical enablers mature, network operators might struggle to operationalise them optimally and economically without sufficient guidance.

Many of above potential or perceived gaps and blind spots are covered through items in our reference framework.

REFERENCES

TM Forum

[1] TM Forum, "IG1339 Autonomous Networks L4 High Value Scenarios v2.4.0", TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1339-autonomous-networks-l4-high-value-scenarios-v2-4-0/. Accessed: Aug. 2, 2026.

[2] TM Forum, "GB1087 Agentic Interaction Security in Telecommunications," Version 3.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/guidebook/gb1087-agentic-interaction-security-in-telecommunications-v3-0-0/. Accessed: Jul. 8, 2026.

[3] TM Forum, "IG1356 Data Architecture for AI-enabled Telecom Operations Whitepaper," Version 3.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1356-data-architecture-for-ai-enabled-telecom-operations-whitepaper-v3-0-0/. Accessed: Jul. 8, 2026.

[4] TM Forum, "Autonomous Networks Framework," IG1218F, Version 2.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/autonomous-networks-framework-v2-0-0-ig1218f/. Accessed: Jul. 8, 2026.

[5] TM Forum, "IG1251C AN Level 4 Target Architecture," Version 2.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1251c-an-level-4-target-architecture-v2-0-0/. Accessed: Jul. 8, 2026.

[6] TM Forum, "IG1251E Multi-Agent Architectures," Version 1.1.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1251e-multi-agent-architectures-v1-1-0/. Accessed: Jul. 8, 2026.

[7] TM Forum, "Autonomous Network Levels Evaluation Methodology," IG1252, Version 1.2.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1252-autonomous-network-levels-evaluation-methodology-v1-2-0/. Accessed: Jul. 8, 2026.

[8] TM Forum, "IG1453 Agent to Agent Protocol for Telecoms," Version 2.1.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1453-agent-to-agent-protocol-for-telecoms-v2-1-0/. Accessed: Jul. 8, 2026.

[9] TM Forum, "IG1463 Security Considerations for Agentic AI in Autonomous Networks", Version v1.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1463-security-considerations-for-agentic-ai-in-autonomous-networks-v1-0-0/. Accessed : Aug. 2, 2026.

[10] TM Forum, "IG1392 Autonomous Networks Levels Assessment and Certification Whitepaper," Version 2.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/ig1392-autonomous-networks-levels-assessment-and-certification-whitepaper-v2-0-0/. Accessed: Jul. 8, 2026.

[11] TM Forum, "IG1412 AI Agent Specification Template," Version 2.0.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/guidebook/ig1412-ai-agent-specification-template-v2-0-0/. Accessed: Jul. 8, 2026.

[12] TM Forum, "Scaling to AN Level 4 to unlock network value," TM Forum Inform. [Online]. Available: https://inform.tmforum.org/research-and-analysis/reports/scaling-to-an-level-4-to-unlock-network-value. Accessed: Jul. 8, 2026.

[13] TM Forum, "Intent Management API," TMF921, Version 5.0, TM Forum Open APIs. [Online]. Available: https://www.tmforum.org/open-digital-architecture/open-apis/intent-management-api-TMF921/v5.0. Accessed: Jul. 8, 2026.

[14] TM Forum, "TR292 TM Forum Intent Ontology (TIO)," Version 3.6.0, TM Forum. [Online]. Available: https://www.tmforum.org/resources/introductory-guide/tr292-tm-forum-intent-ontology-tio-v3-6-0/. Accessed: Jul. 8, 2026.

[15] TM Forum, "AI-Native Blueprint Project," TM Forum. [Online]. Available: https://www.tmforum.org/ai-native-blueprint-project/. Accessed: Jul. 8, 2026.

[16] TM Forum, "Autonomous Networks Project," TM Forum. [Online]. Available:https://www.tmforum.org/autonomous-networks-project/. Accessed: Jul. 8, 2026.

[17] TM Forum, "Business-aware GNN-healing networks," TM Forum Catalyst Project. [Online]. Available: https://www.tmforum.org/catalysts/projects/C26.0.965/businessaware-gnnhealing-networks. Accessed: Jul. 8, 2026.

[18] Google Cloud, TM Forum, and Vodafone, "Self-Optimizing Autonomous Networks: An implementation guide," Jun. 2026. [Online]. Available: https://www.vodafone.com/news/newsroom/technology/vodafone-google-cloud-tm-forum-unveil-framework-for-self-optimising-autonomous-networks.

3GPP

[19] ETSI, "5G; Management and orchestration; Levels of autonomous network," 3GPP TS 28.100, Version 19.0.0, Release 19, ETSI. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/128100_128199/128100/19.00.00_60/ts_128100v190000p.pdf. Accessed: Jul. 8, 2026.

[20] ETSI, "LTE; 5G; Management and orchestration; Intent driven management services for mobile networks," 3GPP TS 28.312, Version 18.4.0, Release 18, ETSI. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/128300_128399/128312/18.04.00_60/ts_128312v180400p.pdf. Accessed: Jul. 8, 2026.

[21] 3GPP, "Study on 6G Use Cases and Service Requirements," Specification No. 22.870, 3GPP. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=4374. Accessed: Jul. 8, 2026.

[22] 3GPP, "Study on 6G Management and Orchestration (Release 20)", 3GPP TR 32.801-01, V0.3.0, June 2026. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=5491. Accessed: Aug 2, 2026.

ETSI

[23] ETSI, "ISG ZSM," ETSI. [Online]. Available: https://www.etsi.org/technical-groups/zsm/. Accessed: Jul. 8, 2026.

[24] ETSI, "ISG ENI," ETSI. [Online]. Available: https://www.etsi.org/technical-groups/eni/. Accessed: Jul. 8, 2026.

[25] Markus Mueck, Luigi Licciardi, "AI and Autonomous Networks: Coordination across ETSI", ETSI, Feb. 11, 2025. [Online]. Available: https://docbox.etsi.org/Workshop/2025/02_AICONFERENCE/SESSION04/ETSI_OCG_AI_AN_MUECK_MARKUS_LICCIARDI_LUIGI.pdf. Accessed : Aug. 2, 2026.

ORAN

[26] O-RAN Software Community, "O-RAN SC Non-RT RIC project: nonrtric master documentation," O-RAN SC. [Online]. Available: https://docs.o-ran-sc.org/projects/o-ran-sc-nonrtric/en/cherry/overview.html. Accessed: Jul. 8, 2026.

[27] O-RAN Alliance, "O-RAN Specifications," O-RAN Alliance. [Online]. Available: https://www.o-ran.org/specifications. Accessed: Jul. 8, 2026.

AI-RAN Alliance

[28] Ai-RAN Alliance, "AI-RAN Alliance Reaches Major Milestone, Showcasing Breakthrough Momentum for AI-Native Networks", 26 Feb 2026. [Online]. Available: https://ai-ran.org/press-releases/mwc-2026-momentum. Accessed: Jul. 14, 2026.

Linux Foundation Agent2Agent (A2A)

[29] Agent2Agent, [Online]. Available: https://a2a-protocol.org/latest/

Linux Foundation AAIF

[30] [Model Context Protocol, "The 2026-07-28 MCP Specification Release Candidate," Model Context Protocol Blog. [Online]. Available: https://blog.modelcontextprotocol.io/posts/2026-07-28-release-candidate/. Accessed: Jul. 8, 2026.

[31] Model Context Protocol, "Specification," Model Context Protocol. [Online]. Available: https://modelcontextprotocol.info/specification/. Accessed: Jul. 8, 2026.

[32] Model Context Protocol, [Online]. Available: https://aaif.io/projects/model-context-protocol/

[33] Agentgateway, [Online]. Available: https://agentgateway.dev/

Linux Foundation AGNTCY

[34] AGNTCY, "Building the Internet of Agents (IoA)," AGNTCY. [Online]. Available: https://agntcy.org/. Accessed: Jul. 8, 2026.

Linux Foundation other projects

[35] CAMARA Project, "CAMARA Project," Linux Foundation Project. [Online]. Available: https://camaraproject.org/. Accessed: Jul. 8, 2026.

[36] Nephio, "Nephio," Linux Foundation Project. [Online]. Available: https://nephio.org/. Accessed: Jul. 8, 2026.

[37] Sylva Project, "Home," Sylva. [Online]. Available: https://sylvaproject.org/home/. Accessed: Jul. 8, 2026.

[38] LF Edge, "Building an Open Source Framework for the Edge," Linux Foundation Project. [Online]. Available: https://lfedge.org/. Accessed: Jul. 8, 2026.

GSMA

[39] GSMA, "GSMA 6G Community - Networks, AI Agents and Agentic Protocols for Telecom Networks Whitepaper" GSMA. [Online]. Available: https://www.gsma.com/solutions-and-impact/technologies/networks/join-the-gsma-6g-community/. Accessed: Jul. 8, 2026.

ITU

[40] ITU-T, "Focus Group on Autonomous Networks (FG-AN)," ITU. [Online]. Available: https://www.itu.int/en/ITU-T/focusgroups/an/Pages/default.aspx. Accessed: Jul. 8, 2026.

IETF

[41] IETF, "Agentic AI Use Cases and Requirements," Internet-Draft draft-agentic-ai-usecases-requirements-00, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-agentic-ai-usecases-requirements. Accessed: Jul. 8, 2026.

[42] IETF, "Security Requirements for AI Agents," Internet-Draft draft-ni-a2a-ai-agent-security-requirements-01, IETF. [Online]. Available: https://www.ietf.org/archive/id/draft-ni-a2a-ai-agent-security-requirements-01.html. Accessed: Jul. 8, 2026.

[43] IETF, "Agentic AI Architectural Principles for Autonomous Computer Networks," Internet-Draft draft-jadoon-nmrg-agentic-ai-autonomous-networks-00, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-jadoon-nmrg-agentic-ai-autonomous-networks/. Accessed: Jul. 8, 2026.

[44] IETF, "AI Agent Discovery and Invocation Protocol," Internet-Draft draft-cui-ai-agent-discovery-invocation-01, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-cui-ai-agent-discovery-invocation/01/. Accessed: Jul. 8, 2026.

[45] IETF, "Agent Name Service (ANS): A Universal Directory for Secure AI Agent Discovery and Interoperability," Internet-Draft draft-narajala-ans-00, IETF. [Online]. Available: https://www.ietf.org/archive/id/draft-narajala-ans-00.html. Accessed: Jul. 8, 2026.

[46] IETF, "Task discovery in agentic networks," Internet-Draft draft-mapmw-task-discovery-00, IETF. [Online]. Available: https://www.ietf.org/archive/id/draft-mapmw-task-discovery-00.html. Accessed: Jul. 8, 2026.

[47] IETF, "DNS-Based Agent Naming (DAN): AIDISCA and AIINDEX Resource Records for AI Agent Discovery," Internet-Draft draft-seethiraju-dawn-dan-00, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-seethiraju-dawn-dan/. Accessed: Jul. 8, 2026.

[48] IETF, "Taxonomy for Agentic AI Use Cases," Internet-Draft draft-scrm-aiproto-usecases-04, IETF. [Online]. Available: https://www.ietf.org/archive/id/draft-scrm-aiproto-usecases-04.html. Accessed: Jul. 8, 2026.

[49] IETF, "Agent Communication Protocols," BOF request bofreq-krishnan-agent-communication-protocols-02, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/bofreq-krishnan-agent-communication-protocols/. Accessed: Jul. 8, 2026.

[50] IETF, "Motivations and Problem Statement of Agentic AI for network management," Internet-Draft draft-hong-nmrg-agenticai-ps-01, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-hong-nmrg-agenticai-ps/. Accessed: Jul. 8, 2026.

[51] IETF, "AI Agent Execution Profile of SCITT," Internet-Draft draft-emirdag-scitt-ai-agent-execution-00, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-emirdag-scitt-ai-agent-execution/. Accessed: Jul. 8, 2026.

[52] IETF, "Agent Identity, Trust and Lifecycle Protocol (AITLP)", Internet-Draft. [Online]. Available: https://datatracker.ietf.org/doc/draft-larsson-aitlp/00/. Accessed : Aug 2, 2026.

IRTF NMRG

[53] IRTF, "Network Management Research Group (NMRG)," IRTF. [Online]. Available: https://www.irtf.org/nmrg.html. Accessed: Jul. 8, 2026.

[54] Y. Cui and M. Xing, "A Framework for LLM Agent-Assisted Network Management with Human-in-the-Loop," Internet-Draft, Jul. 1, 2026, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-cui-nmrg-llm-nm/. Accessed: Jul. 8, 2026.

[55] J. Zhao et al., "AI Agent Architecture for Network Digital Twin," Internet-Draft, Mar. 2, 2026, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-zhao-nmrg-ai-agent-for-ndt/. Accessed: Jul. 8, 2026.

[56] Q. Wu et al., "Network Digital Twin and Agentic AI based Architecture for AI driven Network Operations," Internet-Draft, May 21, 2026, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-wmz-nmrg-agent-ndt-arch/. Accessed: Jul. 8, 2026.

[57] M. A. Jadoon et al., "Agentic AI Architectural Principles for Autonomous Computer Networks," Internet-Draft, Mar. 2, 2026, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-jadoon-nmrg-agentic-ai-autonomous-networks/. Accessed: Jul. 8, 2026.

[58] J. Francois et al., "Research Challenges in Coupling Artificial Intelligence and Network Management," Internet-Draft, Jul. 6, 2026, IETF Datatracker. [Online]. Available: https://datatracker.ietf.org/doc/draft-irtf-nmrg-ai-challenges/. Accessed: Jul. 8, 2026.

BBF

[59] Broadband Forum, "AI in Broadband: From vision to blueprint", 3 July 2026 [Online]. Available: https://www.broadband-forum.org/blog/ai-in-broadband-from-vision-to-blueprint/. Accessed: Jul. 13, 2026.

[60] Broadband Forum, "MR-529 AI in Broadband Networks", Issue: 1, May 2026, [Online]. Available: https://www.broadband-forum.org/pdfs/mr-529-1-0-0.pdf. Accessed: Jul. 13, 2026.

W3C

[61] W3C, "AI Agent Protocol Community Group," W3C. [Online]. Available: https://www.w3.org/community/agentprotocol/. Accessed: Jul. 8, 2026.

[62] W3C, "AI KR (Artificial Intelligence Knowledge Representation) Community Group," W3C. [Online]. Available: https://www.w3.org/community/aikr/. Accessed: Jul. 8, 2026.

[63] W3C, "Semantic Agent Communication Community Group," W3C. [Online]. Available: https://www.w3.org/community/s-agent-comm/. Accessed: Jul. 8, 2026.

[64] W3C, "Agent Declaration and Assurance Community Group," W3C. [Online]. Available: https://www.w3.org/community/adacg/. Accessed: Jul. 8, 2026.

[65] W3C, "Semantic Web Interest Group," W3C. [Online]. Available: https://www.w3.org/groups/ig/semweb/. Accessed: Jul. 8, 2026.

[66] W3C, "Credentials Community Group," W3C. [Online]. Available: https://www.w3.org/community/credentials/. Accessed: Jul. 8, 2026.

[67] W3C, "Decentralised Identifiers (DIDs) v1.0," W3C Recommendation. [Online]. Available: https://www.w3.org/TR/did-1.0/. Accessed: Jul. 8, 2026.

[68] W3C, "Verifiable Credentials Data Model v2.0," W3C Recommendation. [Online]. Available: https://www.w3.org/TR/vc-data-model-2.0/. Accessed: Jul. 8, 2026.

OPENID

[69] OpenID Foundation, "Artificial Intelligence Identity Management Community Group," OpenID Foundation. [Online]. Available: https://openid.net/cg/artificial-intelligence-identity-management-community-group/. Accessed: Jul. 8, 2026.

Cloud Security Alliance

[70] Cloud Security Alliance, "Agentic AI Governance: NIST Standards for Autonomous Systems", [Online]. Available: https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/03/governance-nist-ai-agent-standards-agentic-governance-v1-csa-styled.pdf. Accessed: Jul. 8, 2026.

NIST

[71] NIST, "AI Agent Standards Initiative," National Institute of Standards and Technology. [Online]. Available: https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative. Accessed: Jul. 8, 2026.

NGMN Alliance

[72] NGMN Alliance, "Cloud-native next chapter - Agentic AI-based operating models". Version 1.0. Mar. 24, 2026 [Online]. Available: 2603_NGMN_RTD_-Agentic-AI-Based-Operating-Models_V1.0.pdf. Accessed: Aug. 4, 2026.

Tokenomics

[73] GitHub, "AI Credits Specification," GitHub Agentic Workflows. [Online]. Available: https://github.github.com/gh-aw/specs/ai-credits-specification/. Accessed: Jul. 8, 2026.

[74] OWASP Foundation, "OWASP Agent Observability Standard," OWASP Foundation. [Online]. Available: https://owasp.org/www-project-agent-observability-standard-2/. Accessed: Jul. 8, 2026.

[75] Stanford Digital Economy Lab, "How Do AI Agents Spend Your Money? Analyzing and Predicting Token Consumption in Agentic Coding Tasks," Stanford Digital Economy Lab. [Online]. Available: https://digitaleconomy.stanford.edu/publication/how-do-ai-agents-spend-your-money-analyzing-and-predicting-token-consumption-in-agentic-coding-tasks/. Accessed: Jul. 8, 2026.

[76] Gartner, "Gartner Predicts AI Coding Costs Will Surpass Average Developer's Salary by 2028 as Token Consumption Surges," Jun. 24, 2026, Gartner Newsroom. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2026-06-24-gartner-predicts-ai-coding-costs-will-surpass-average-developer-salary-by-2028-as-token-consumption-surges. Accessed: Jul. 8, 2026.

[77] Mavik Labs, "LLM Cost Optimisation in 2026: Routing, Caching, and Batching," Mavik Labs. [Online]. Available: https://www.maviklabs.com/blog/llm-cost-optimisation-2026. Accessed: Jul. 8, 2026.

Agent interpretability

[78] Anthropic, "Natural Language Autoencoders: Turning Claude's thoughts into text," Anthropic Research. [Online]. Available: https://www.anthropic.com/research/natural-language-autoencoders. Accessed: Jul. 8, 2026.

Agent trustworthiness

[79] Anthropic, "Trustworthy agents in practice," Anthropic Research. [Online]. Available: https://www.anthropic.com/research/trustworthy-agents. Accessed: Jul. 8, 2026.

Evaluation of agents

[80] Anthropic, "Demystifying evals for AI agents," Anthropic Engineering. [Online]. Available: https://www.anthropic.com/engineering/demystifying-evals-for-ai-agents. Accessed: Jul. 8, 2026.

[81] TeleCom-Bench authors, "TeleCom-Bench: How Far Are Large Language Models from Industrial Telecommunication Applications?," arXiv. [Online]. Available: https://arxiv.org/html/2605.18025v1. Accessed: Jul. 8, 2026.

Measuring agent autonomy

[82] Anthropic, "Measuring AI agent autonomy in practice," Anthropic Research. [Online]. Available: https://www.anthropic.com/research/measuring-agent-autonomy. Accessed: Jul. 8, 2026.

Governance frameworks for agentic AI

[83] IMDA, "Model AI Governance Framework for Agentic AI," Infocomm Media Development

Authority. [Online]. Available: https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf. Accessed: Jul. 8, 2026.

Digital twins for networks

[84] VIAVI Solutions, "Unlock the Power of Digital Twins: AI-Driven Network Optimisation & Automation," VIAVI Solutions. [Online]. Available: https://images.comms.viavisolutions.com/Web/Viavi/%7B9c7e9350-4d7a-40a5-b7c9-959fa13b0931%7D_digitaltwin-wp-wir-nse-ae.pdf. Accessed: Jul. 8, 2026.

White papers and frameworks

[85] ZTE, "ZTE Autonomous Networks White Paper 2026," 2026, ZTE. [Online]. Available: https://www.zte.com.cn/content/dam/zte-site/res-www-zte-com-cn/mediares/zte/service/ZTE%20Autonomous%20Networks%20White%20Paper%202026%20\(EN\).pdf. Accessed: Jul. 8, 2026.

[86] Vodafone, "Vodafone, Google Cloud and TM Forum Unveil Framework for Self-Optimising Autonomous Networks," Vodafone Newsroom. [Online]. Available: https://www.vodafone.com/news/newsroom/technology/vodafone-google-cloud-tm-forum-unveil-framework-for-self-optimising-autonomous-networks. Accessed: Jul. 8, 2026.

Other relevant documents

[87] Anthropic, "Building Effective AI Agents," Anthropic Engineering. [Online]. Available: https://www.anthropic.com/engineering/building-effective-agents. Accessed: Jul. 8, 2026.

[88] Spherical Cow Consulting, "Agentic AI in the Open Standards Community: Standards Work or Just Hype?," Aug. 12, 2025, Spherical Cow Consulting. [Online]. Available: https://sphericalcowconsulting.com/2025/08/12/agentic-ai-and-open-standards/. Accessed: Jul. 8, 2026.

[89] R. K. Sharma, "Criticality-Based Guard Rail Validation for AI Agent Decisions in Autonomous Telecom Networks," Jul. 2, 2026, arXiv. [Online]. Available: https://arxiv.org/html/2607.02210v1. Accessed: Jul. 8, 2026.

[90] Z. Nezami et al., "Safety and Risk Pathways in Cooperative Generative Multi-Agent Systems: A Telecom Perspective," Nov. 21, 2025, arXiv. [Online]. Available: https://arxiv.org/pdf/2511.17730. Accessed: Jul. 8, 2026.

[91] Maxime Elkael†, Salvatore D'Oro, Leonardo Bonati, Michele Polese, Yunseong Lee, Koichiro Furueda, Tommaso Melodia, "AgentRAN: AnAgentic AI Architecture for Autonomous Control of Open 6G Network", 2 Feb 2026. [Online]. Available: https://arxiv.org/pdf/2508.17778

[92] Yunhao Hu, Xinchen Lyu, Chenshan Ren, Keda Chen, Qimei Cui, and Xiaofeng Tap, "Reflection-Driven Self-Optimisation 6G Agentic AI RAN via Simulation-in-the-Loop Workflows", 21 Apr 2026. [Online]. Available: https://arxiv.org/pdf/2512.20640

FIGURES

Figure 1:
Illustrative example
for a fault management workflow 10

Figure 2:
Illustrative example
for a service assurance workflow 12

Figure 3:
Illustrative example
for a RAN optimisation workflow 14

Figure 4:
Overall framework of core
components – Network Automation
with Agentic AI 16

ACKNOWLEDGEMENTS

China Mobile, Lingli Deng	Project Lead
Vodafone, Guenter Klas	Project Co-Lead
NGMN, Johann Reindl	NGMN Programme Manager (external)

Special thanks to the following organisations and experts for their contributions:

China Mobile, Lei Huang	**Ericsson,** Dinand Roeland
China Mobile, Sheng Gao	**Nokia,** Kurt Pynaert,
Liberty Global, Assen Golaup	**Nokia,** Renata Silva
P3, Sayer Khandker	**Nokia,** Marko Salmimaa
Vodafone, Simone Mangiante	**Nokia,** Nils Kleemann
Vodafone, Joerg Sternagel	**Dell,** Neeraj Sharma
Ericsson, Ignacio Mas	**ZTE,** Wang Yuqi

The project team also wishes to thank the experts who reviewed an interim draft of this publication: Kurt Pynaert (Nokia), Richard Kilmurray (Vodafone), Andrey Silva (Ericsson), Mateus Santos (Ericsson), Sergio Benco (Telecom Italia).

NEXT GENERATION MOBILE NETWORKS ALLIANCE

NGMN - Next Generation Mobile Networks Alliance - is a global, operator-driven organisation established by leading international mobile network operators (MNOs). As a global alliance of operators, vendors, and academia, NGMN provides industry guidance to enable innovative, sustainable and affordable next-generation mobile network infrastructure.

Key focus areas include Mastering the Route to Disaggregation, Green Future Networks, and 6G, while supporting the full implementation of 5G. NGMN drives global alignment of technology standards, fosters collaboration with industry organisations and ensures efficient, project-driven processes to address the evolving demands of the telecommunications ecosystem.

VISION

The vision of NGMN is to provide impactful industry guidance to achieve innovative, sustainable and affordable mobile telecommunication services to meet the requirements of operators and address the demands and expectations of end users. Key focus areas include Mastering the Route to Disaggregation, Green Future Networks and 6G, while supporting the full implementation of 5G.

MISSION

The mission of NGMN is:

- To evaluate and drive technology evolution towards the three **Strategic Focus Topics:**
  - **Mastering to the Route to Disaggregation:**

    Leading in the development of open, disaggregated, virtualised and cloud-native solutions, moving towards agentic AI-based operating models.
  - **Green Future Networks:**

    Building sustainable and environmentally conscious solutions.
  - **6G:**

    Providing guidance and key requirements for design considerations and network architecture evolution.
- To define precise functional and non-functional requirements for the next generation of mobile networks.
- To provide guidance to equipment developers, standardisation bodies, and collaborative partners, leading to the implementation of a cost-effective network evolution.
- To serve as a platform for information exchange within the industry, addressing urgent concerns, sharing experiences, and learning from technological challenges.
- To identify and eliminate obstacles hindering the successful implementation of appealing mobile services.