ME MOBILE EUROPE | INDEPENDENT RESEARCH

Agentic AI

Achieving autonomous network operations

SUE TABBITT
BUSINESS & TECH JOURNALIST

APRIL 2026

With thanks to our sponsor CELFOCUS

Agentic AI

Achieving autonomous network operations

MEET THE AUTHOR

Sue Tabbitt

Sue has been a technology and business journalist for 35 years, She started on *Computergram International* (which became *Computerwire*), between 1990-1992. After 10 years on tech trade magazines in London, including as editor of *The VAR* magazine and managing editor of *Network Reseller Magazine*, she turned freelance and moved to France. Since 2001, Sue has worked as a tech journalist for trade publications, web sites and national newspapers, as a prolific corporate copywriter and report writer for the Management Consultancies Agency and what was then Source Consulting. She ghost wrote Joined up systems, published by Hodder & Stoughton.Recent reports covered: AI monetisation by telcos; smart cities in Singapore; a history of quantum technology: and the future of pharma supply chains. She writes extensively about tech in life sciences, healthcare and other sectors. Sue works remotely from West Wales.

Editor:
Annie Turner
anniet@mobileeurope.co.uk

Group Commercial Director:
Fidi Neophytou
fidin@mobileeurope.co.uk
Tel +44 7741 911 302

Director of Business Development:
Simon La Thangue,
simonlt@mobileeurope.co.uk
Tel+44 7776 321 644

Marketing Director:
Marcin Piechowski
marcinp@sjpbusinessmedia.com

Design and Production:
Rob Castles
robc@sjpbusinessmedia.com

April 2026
This document is published by
SJP Business Media
37 Lombard Street
London, EC3V 9BQ
Tel +44 (0) 207 062 2526

CONTENTS

Executive summary

Agentic AI will see orchestrated AI systems reason, plan and act autonomously across complex, multi-step workflows, unlike previous forms of AI, which respond to commands or analyse data, but cannot reason or initiate actions. Agentic AI is viewed as the mechanism to close the 'autonomy chasm' in telecoms by bringing an AI-driven, automated approach to managing network operations (ANOps).

The 'autonomy gap' is the perceived difficulty of moving from TM Forum's Level 3 (conditional autonomy) to Level 4 (high autonomy) in its Autonomous Network (AN) framework – see graphic on page 5. Why is bridging the gap important? The Forum's modelling suggests AN maturity could deliver a 55% reduction in operations and maintenance costs, and a 71% rise in customer satisfaction, which would be huge gains for operators facing pressures including:

- Network complexity has grown exponentially with 5G Standalone deployments, network slicing, open RAN and other RAN architectures, and cloud-native infrastructure.
- Traditional rule-based automation cannot scale to manage dynamic, software-defined environments.
- Enterprise customers want telco-grade services delivered with cloud-like simplicity and flexibility, with real-time visibility into performance and proactive resolution of issues.

DATA

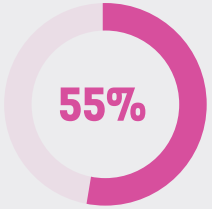

Reduction AN maturity could deliver in operations and maintenance costs

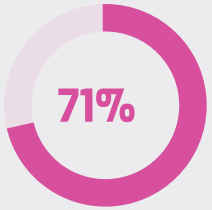

Potential rise in customer satisfaction AN maturity could offer

Source: TM Forum

Momentum to increase automation is building in the industry although two-thirds of the value from network automation has been captured through conventional approaches, according to Charlotte Patrick Research. Further, the economics of full autonomy are far from guaranteed: Gartner expects that more than 40% of agentic AI projects across all enterprise markets will be cancelled by the end of 2027.

Agentic AI and autonomy: CSPs set out their strategies, published by TM Forum last autumn, found that 73% of operators are expecting to run pilots at least in the next two years.

In telecoms, the players making the most progress with ANOps share a common discipline. They start with narrow, well-defined use cases, where return on investment (ROI) is demonstrable and the risk contained. They begin governance for 'AgentOps' on day one, define measurement frameworks before deployment and address organisational change.

This report explores how and how soon the promise of AgentOps within ANOps will it be realised, where investments will pay off, the challenges including economics, what critical success factors are likely to be and recommendation for next steps.

Chapter 1 looks at why ANOps are so important; Chapter 2 digs into how agentic AI bring entirely new dimensions into telecoms operations. Chapter 3 explores how operators are experimenting with and deploying agentic AI in ANOps. The Conclusion makes a series of recommendations for next steps and success.

The report includes use cases and insights from operators including BT, BT Business, Deutsche Telekom, MEO, Orange, Proximus Ada, Swisscom, Telekom Srbija, Telefónica, Telenor, Telstra and Vodafone, plus Google Cloud.

CHAPTER 1

Why do we need AI-enabled autonomous network operations (ANOps)?

Telecoms has long focused on the promise and delivery of autonomous networks (AN) for tasks like rebalancing traffic or healing a RAN fault, and a multitude of others within the network. TM Forum's AN framework is widely recognised and used across the industry (see graphic, right). AI-enabled autonomous network operations (ANOps) cover AN capabilities, but extend AI's reach – through orchestrated, intelligent and coordinated reasoning – to service design and provisioning, customer care and billing, field service dispatch, fraud detection, enterprise SLA management and back-office processes.

Consequently, ANOps hold out great promise to operators as they face relentless pressure to reduce opex while embracing new tech and improving service quality. Manual processes obstruct these goals. Industry analyst Charlotte Patrick models the financial impact of increased autonomy and calculates that two-thirds of automation's potential value has already been captured within the network itself.

She says, "We've done the basic automations, harvested the low-hanging fruit. If we want to do the last piece and get towards autonomy, we probably do need to do agentic but that only gives you a third of that total value to play with."

The last third is hard

That last third might cost more to achieve than it saves in some scenarios, she warns. Once network operations centres (NOCs) have been reduced to minimal staffing, typically to around 12 people at a

QUOTE

"We've done the basic automations, harvested the low-hanging fruit. If we want to do the last piece and get towards autonomy, we probably do need to do agentic but that only gives you a third of that total value to play with"

Charlotte Patrick, Industry Analyst

Autonomous network levels

Source: TM Forum, Exploring the evolution from Level 0 to Level 5, published December 2021

cost of around £400,000 (about €537,000) annually, achieving full autonomy could cost £2 million or more in technology and ongoing operations. Patrick states achieving a "dark" or unstaffed NOC by 2030 is "never going to happen," although it is often cited as a strategic imperative for operators.

Afnan Ahmed, Director of Technology Strategy and Architecture at Telenor, expands on this, saying, "If the answer [to why we're doing automation] is just efficiency, I think we will struggle." In his view, the business case must span multiple dimensions – including improved customer experience, protecting revenue streams and enabling new ones – not simply reducing headcount. He was speaking as a panellist at **Mobile Europe**'s recent *The Briefing, Achieving end-to-end automation*. Watch on playback.

Simplification before automation

5G Standalone (5G SA), network slicing, open RAN – with AI-RAN emerging rapidly – cloud-native architectures and private networks have dramatically increased operational complexity. Ahmed recalls a conversation with a peer at Deutsche Telekom who said they stopped counting operational software systems once they reached 1,000.

"Homegrown overlays" compounds the complication, he notes, adding that rationalising tool sprawl while increasing autonomy requires sophisticated orchestration. "There is no way that anyone can upgrade the tool suite we have from the legacy; that needs to be sunset in a good way," Ahmed says.

In this interview, Mark Düsener, CTIO at Swisscom, explains the operator will have phased out 70% of its network applications and platforms – cutting 15 IP networks down to one – by the end of 2026 to avoid layering AI capabilities on top of legacy technology.

DATA

Swisscom's approach to simplification

15▸1

Reducing 15 IP networks to 1

Decomposing OSS into 30 autonomous operational domains

Source: Interview with Mark Düsener, CTIO at Swisscom

Domain-based approaches

In the same panel as Ahmed, José Palma, Head of Network Strategy at the Portuguese operator MEO, says that with physical, virtualised and containerised infrastructure layers running simultaneously, "it's hard to automate a service across all these platforms. We really need to simplify before we automate". He advocates starting with single-domain use cases rather than attempting end-to-end automation prematurely.

Natali Delić, Chief Strategy and Digital Officer at Telekom Srbija, who also spoke on the panel, said it had moved from a use-case-by-use-case approach to a domain-based governance model. The model evaluates how initiatives contribute collectively to business outcomes, not individual merit.

Swisscom too has taken a domain-based approach. Led by Düsener, the operator decomposed its OSS into 30 autonomous operational domains, based on TM Forum's Open Digital Architecture (ODA) framework. Each component is decoupled so that a dedicated team can work on it without affecting others or having to wait for other teams to complete their part.

Cross-domain complexity

A domain-based approach to autonomy makes sense for the reasons explained above but ultimately operators will need visibility across 5G SA, slicing and cloud-native architectures as they add new service layers. For instance, Patrick observes that for network slicing, operators need "to view transport, core and RAN metrics in a single dashboard for network slices; also, more requirements to provide SLA and KPI insights to slice owners."

Traditional, rules-based automation is not designed to manage dynamic, software-defined environments with numerous variables and interdependencies, but agentic architectures can traverse traditional silos through specialised agents that communicate and cooperate.

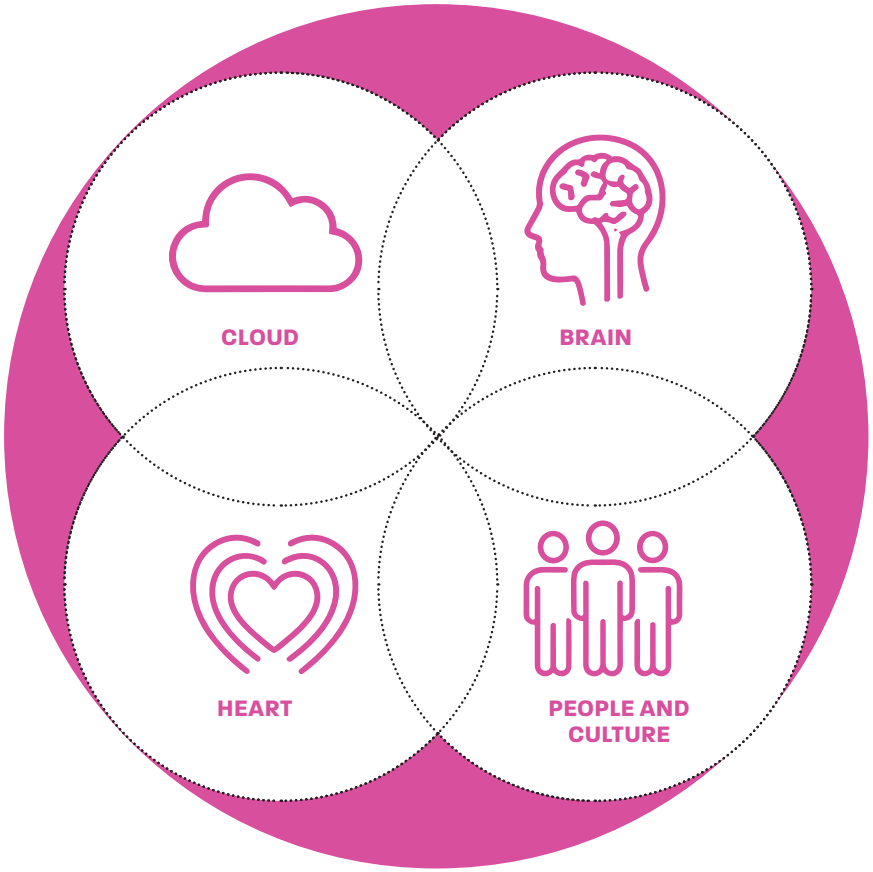

The 4 interlocking workstreams in Telefónica's AN programme, which must develop in parallel to maintain pace

The layered approach to autonomy

Telefónica launched its AN programme in 2021 and maps it onto TM Forum's AN framework. The programme is structured around four interlocking workstreams, according to Nilmar Seccomandi David, Director of AN and infrastructure.

They are:

- An open, disaggregated cloud-based network architecture.
- The 'brain' creates the platform for automated intelligent decision-making comprising a federated data architecture, an orchestration model and modernised OSS.
- The 'heart' involves rethinking and redesigning end-to-end process lifecycles.
- People and culture – redefining roles, introducing new technology and fostering collaboration across silos.

Telefónica's Head of AI, Carolina García, explained the programme's multiple workstreams must advance in parallel to maintain pace, speaking as a panellist at *The Briefing, Becoming an AI-native telco* in January (watch on playback). Country-level AI Centres of Excellence provide shared platforms and capabilities for GenAI, agentic AI and retrieval-augmented generation (RAG) architectures.

The latter is an AI framework which combines semantic search with generative AI to improve AI accuracy. Telefónica is testing Model Context Protocol (MCP) and the agent-to-agent (A2A) protocol to connect agents with data, systems and each other, all of which we explore in Chapter 2.

Orange's multi-domain approach

For Orange, the operating model, rather than specific tools or processes, provides the fundamental enabler in achieving reliable

autonomy. “You can have the best tools in the world, but if the operating model is not making things connect and work properly, you will not achieve any autonomy,” said Alexis Koalla, Director of Operations Strategy and Transformation, during his keynote at *The Briefing, From autonomous networks to autonomous ops* (watch here on playback).

Orange uses TM Forum’s AN framework not as a target but as a diagnostic tool. The operator’s AI-powered operating model, Atom, a work in progress, is designed to embed AI and automation throughout operations in parallel with upskilling staff and redesigning processes.

Koalla notes that the picture is uneven for Orange’s change management domain across the operator’s 26 affiliates. For instance, deployment and release sit at around Level 3 with Level 4 within reach using existing automation. Monitoring and fault management remain between Levels 2 and 3, held back by data ingestion and pipeline gaps.

Hence Orange is working towards a single common data model, prioritising key data and making it accessible to whoever needs it. Each team maintains business-specific models built on shared-access patterns. The operator has defined clear responsibilities with partitions between Orange, its partners and vendors, and deploys strong governance to pre-empt issues.

Enterprises' expectations rise

As enterprise customers begin deploying advanced AI capabilities at scale to improve their own efficiency, the networks that serve them will need to meet their demands. Operators with B2B customers are already engineering for this. For example, BT Business’ services underpin critical national infrastructure in the UK, including emergency services. It is building Layer 2-3 network fabric (see graphic on page 5) for three new traffic patterns it will need to cater for.

QUOTE

“You can have the best tools in the world, but if the operating model is not making things connect and work properly, you will not achieve any autonomy”

Alexis Koalla, Director of Operations Strategy and Transformation at Orange

Colin Bannon, CTO at BT Business, describes the new traffic patterns as:

- agent-to-agent communications which are continuous, at machine speed;
- post-training reasoning flows – when large language models (LLMs) process complex queries they generate very large, lengthy data transfers requiring consistently high throughput without packet loss; and
- asymmetric upstream flows, such as generated by smart glasses and augmented reality (AR) applications, reverse traditional traffic patterns in mobile network (where downloads dominate) because they generate high volumes of upload traffic.

Bannon was speaking in a recent interview as part of *The Briefing entitled AI, efficiency and profit* (watch here on playback).

We look at agentic AI’s attributes and prospects in the next chapter.

CHAPTER 2

How is agentic AI different from other kinds of AI and how can telcos prepare for it?

GenAI has had a big impact on handling content and knowledge but it cannot deliver autonomous operations. GenAI's capabilities are drawn from large language models (LLMs) which by design are stateless, meaning each prompt is a new interaction, with no memory of previous messages. GenAI cannot initiate actions, are at risk of hallucinating and operate in isolation, away from live operational systems. Agentic AI addresses these limitations by combining reasoning and execution.

Gartner warns that much of what vendors describe as AI agents are in fact relabelled AI-enabled assistants, a practice it calls agent-washing. This could confuse the picture as operators strive to measure their progress against the market. The essential components of an agent are shown in the graphic right.

Leading edge evolution

Gartner positions agentic AI at the leading edge of the AI evolution: traditional (analytical) AI and LLM-based systems act when prompted by human inputs. They don't initiate actions.

Agentic AI can: receive and act on high-level goals; decompose goals into steps; select appropriate tools; take action; and adapt based on outcomes. Agentic systems can learn from their environment, make decisions and perform tasks independently.

Across all sectors, not just telecoms, Gartner predicts that:

- By 2028, 33% of enterprise software applications will include agentic AI and that at least 15% of daily work decisions will be made autonomously

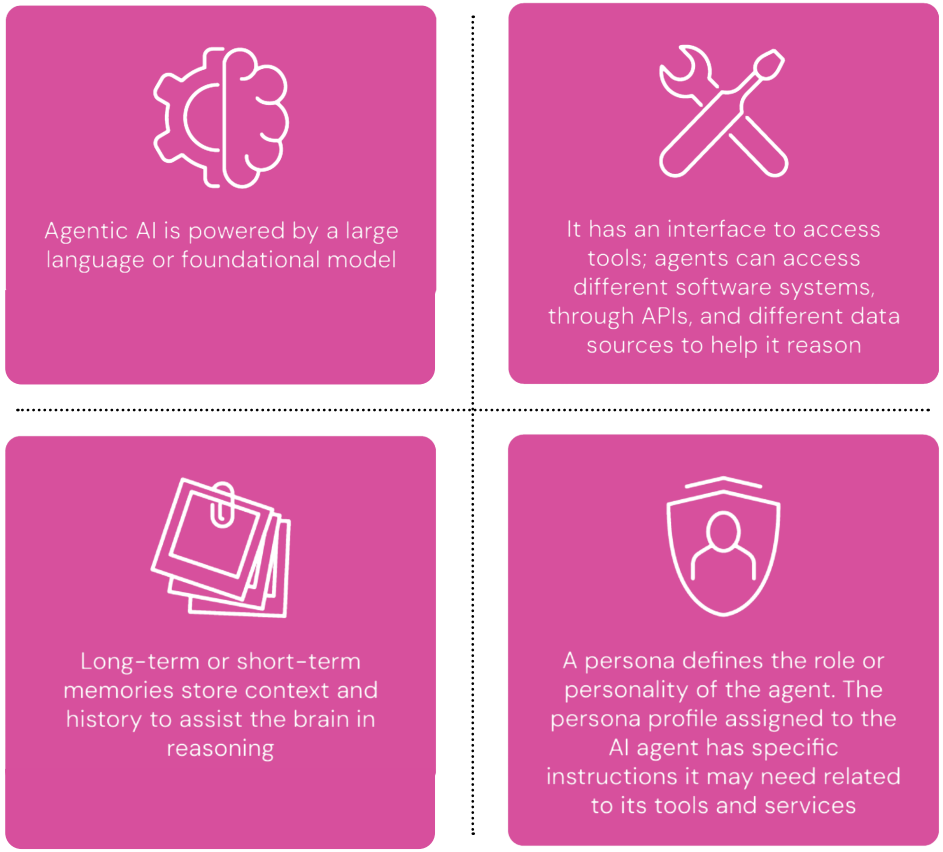

These are the essential components of agentic AI

Source: *Agentic AI and autonomy: CSPs set out their strategies*, published by TM Forum, September 2025

- By 2029, 70% of enterprises will deploy agentic AI as part of IT infrastructure operations, up from less than 5% in 2025
- By 2035, at best agentic AI could drive about 30% of software revenue for enterprise applications, surpassing $450 billion.

Gartner also predicts that more than 40% of agentic AI projects will be cancelled by end of 2027 due to escalating costs, unclear business value or inadequate risk controls.

Agentic AI and autonomy: CSPs set out their strategies, published by TM Forum last autumn, found that 73% of operators are expecting to run pilots at least in the next two years (see pie chart). It surveyed more than 100 respondents from 68 operators around the world.

Agency, reasoning and execution

Philippe Ensarguet, Orange's VP of Software Engineering, describes agents as "autonomous, interactive, goal-oriented, reactive and proactive, task-specific." Where "LLMs think" (provide reasoning), "RAG knows" (retrieval-augmented generation supplies context and knowledge), and "agents do" (execute actions autonomously), he explains in his keynote conversation for *The Briefing* on *Becoming an AI-native telco* (watch on playback).

Traditional automation executes predefined workflows: agentic capability reasons what workflow is needed. Agentic AI adapts its approach, based on context and can handle situations without being explicitly programmed to do so. This adaptability means agents can work across organisational silos, unlike all previous automation technologies.

Mapping where value will come

Analyst Charlotte Patrick tracks the evolution of agentic AI to develop a pragmatic framework for understanding where and when its value will come in telecoms. Rather than focusing on distant visions of full autonomy, her model maps realistic near-term progress across three

QUOTE

"...agents [are] 'autonomous, interactive, goal-oriented, reactive and proactive, task-specific'. Where 'LLMs think' (provide reasoning), 'RAG knows' (retrieval-augmented generation supplies context and knowledge), and 'agents do' (execute actions autonomously)"

Philippe Ensarguet, VP of Software Engineering, Orange

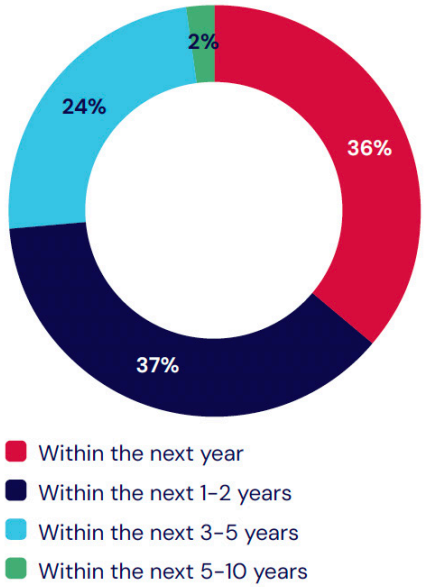

When do you think you will be adopting AI pilots / agents widely across your business?

Source: Source: Agentic AI and autonomy: CSPs set out their strategies, published by TM Forum, September 2025

stages (see graphic on next page). Her staging is designed to help operators set realistic expectations around autonomous operations enabled by agentic AI and discourage over-investment in approaches that lack solid foundations.

Most operators who are underway with agentic AI are making progress in Stage 1, deploying copilot agents supported by simple agent hierarchies. A copilot agent interprets natural language requests, delegates to specialised sub-agents and returns validated results. Patrick expects broad roll-out in relatively low-risk, data-rich domains like service assurance within one to two years, given that

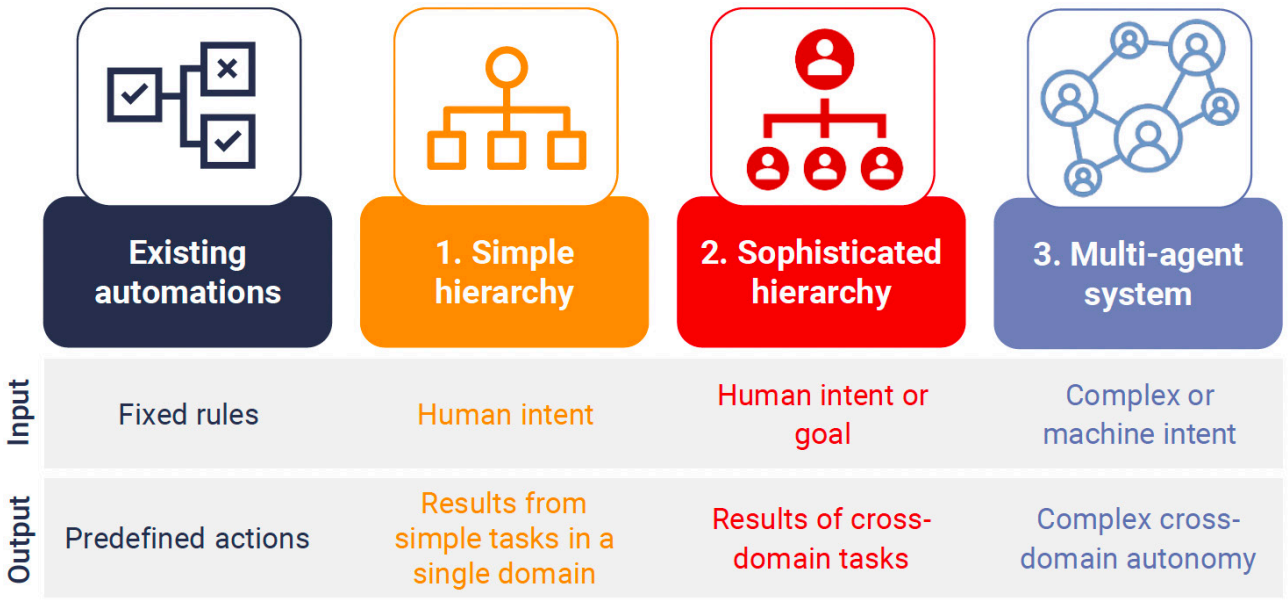

The three-stage journey towards complex agentic systems

Source: Charlotte Patrick Research's *three-stage agentic AI evolution framework, from AI and Agents in Next-generation Assurance* executive briefing by Charlotte Patrick, Associate Analyst, STL Partners

assurance allows humans to remain in the loop and mistakes carry manageable consequences.

Patrick projects Stage 2 will take two to three years to deploy at scale. It will see operators move to more sophisticated hierarchies as agents gain the ability to reason over structured knowledge graphs, making them more reliable and allowing them to expand into faster operational processes (see Knowledge and context on page 14 below).

A knowledge graph is a data structure that connects real-world entities – such as people, places or concepts – and their relationships in a network, typically using nodes (entities) and edges (relationships). As yet, no approach has been found to enabling an LLM to reason correctly and reliably over a knowledge graph at production scale.

Stage 3 represents true multi-agent systems with distributed intelligence and is a much longer term prospect, at least five years away. In such scenarios, agents would become autonomous peers – negotiating with each other rather than following centralised orchestration – closer to TM Forum's Level 5 network autonomy (see graphic on page 5).

As yet there is no resolution to conflicts between agents pursuing competing goals. Patrick notes that situations where fully distributed, multi-agent intelligence is genuinely required may be less common than hype suggests.

In the more immediate future, critical foundational work is necessary before agentic AI can be applied, especially given its promise of operational autonomy.

Emerging protocols

At this nascent stage, there are technical facilitators to help operators progress their agentic AI ambitions, including protocols, most notably the model context protocol (MCP) and agent-to-agent (A2A) protocol.

The MCP provides vertical integration between agents and tools, resources and prompts. Note that those designed for public contexts and consumer applications are unlikely to offer enterprise-grade security without reinforcement.

The A2A protocol enables agent-to-agent collaboration through synchronous and asynchronous communication patterns, meaning agents can request information or actions from other agents and coordinate responses.

Even with these facilities, Orange's Ensarguet warns that considerable complexity remains around deployment, operation, validation and testing. TM Forum's Project Foundation seeks to help with this. It is building the industry's first AI-Native Open Digital Architecture (ODA) Canvas Sandbox which is a Kubernetes-orchestrated environment, that is, an open-source platform where telcos, hyperscalers and partners can co-develop, integrate and test interoperable AI agents aligned with TM Forum's AI-Native Blueprint.

DATA

Effort in agentic AI

Redesigning processes and "hard questions about data"

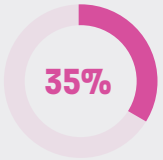

Identifying or building APIs

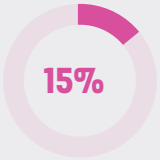

Developing the agent itself

Source Mark Düsener, CTIO, Swisscom

Infrastructure for building

Agent development kits (ADKs) are intended to act as foundational infrastructure for building, deploying and managing agents – typically covering orchestration, integrating tools, managing memory, guardrails and observability. In 2025, a number of agentic AI toolkits for telecoms came onto the market, with technology vendors releasing telco-specific frameworks, reference architectures and agent-building platforms.

They variously targeted network operations, OSS/BSS automation and customer service, but none provides quick fixes for data quality and integration challenges, highlighted succinctly by Swisscom's CTIO, Mark Düsener.

He says understanding the current state, identifying or creating the necessary APIs account for about 35% of the effort required; 50% goes into redesigning processes and asking "hard questions about data". Only 15% of the work is developing the agent itself. "The last, small part is building the agent, yet that's where the public focus mostly is," he notes.

The critical importance of AgentOps

Patrick warns that operators risk repeating the mistakes of early robotic process automation (RPA) deployments if they do not get the foundations right before throwing everything at evolving AI options.

In those early automation scenarios, organisations often allowed teams to deploy the technology everywhere, creating fragmentation and management nightmares. They then had to retrofit governance, security and lifecycle management, which typically was neither efficient nor effective. The same risk exists with agents today, Patrick says, without proper AgentOps – that is, operational frameworks

analogous to DevOps for software or MLOps for machine learning models.

AgentOps encompass governance which:

- determines who can create agents and what permissions they have
- lifecycle management, that is from design to deployment, monitoring, updating and retirement
- security covering access controls, audit logging and prevention of attacks
- coordination controls how agents discover and interact with each other.

Unless AgentOps are established from the start, the tactical proliferation of pilots could create technical debt, gaps in security and coordination failures that become exponentially harder to fix as deployments scale.

Knowledge and context

Agentic AI's success relies on the range and value of the information it has to work with, which is where knowledge graphs, ontologies and curated documentation come into play.

Mark Sanders, Head of Network Transformation at Telstra argues that the industry needs "a shared language, an ontology of how we can model knowledge for a telco business" that goes far beyond network topology. "It needs to have business processes coded in, our engineering design limits, our policies and even regulatory frameworks," he told TM Forum's Inform in February 2026. "Those things need to be coded into the knowledge behind it."

In the same interview, Dennis Sehalic, Senior Solution Architect at Telenor Sweden, links this to TM Forum's AN framework, noting that achieving Level 4 means "orders of magnitude more automations and more meshed automations" which makes it essential to classify and control individual parts.

QUOTE

"[Industry needs] a shared language, an ontology of how we can model knowledge for a telco business [beyond network topology]. It needs to have business processes coded in, our engineering design limits, our policies and even regulatory frameworks. Those things need to be coded into the knowledge behind it"

Source: Mark Sanders, Head of Network Transformation at Telstra

With contextual knowledge, an agent can reason holistically rather than treating a specific situation or incident. For example, when diagnosing service degradation, a system can immediately understand:

- which customers are affected
- what their contract terms are
- which network elements serve them
- what changes were recently made, similar past incidents and their resolutions – synthesising this context
- then recommend and autonomously execute appropriate action.

New operational discipline

Not surprisingly, progress with all of this has been slow: many telcos are yet to start building more than their first, small knowledge assets, according to Patrick. Creating and maintaining high-quality knowledge will require ongoing curation, validation and updates – essentially establishing a new operational discipline.

This is unglamorous, methodical work but without it agents would have to operate with incomplete understanding, limiting their effectiveness.

Despite these challenges and so much work in progress, ambitious if cautious operators are keen to seize advantage through agentic AI, as we see in the next chapter.

CHAPTER 3

Case studies: putting agentic AI to work

Many communications service providers the world over feel acutely that they 'missed out' on cloud computing; they failed to build capabilities to become players and are 'only' users. With AI, they are determined to lead and gain first-adopter advantage, according to Mark Newman, Chief Analyst, TM Forum. He is author of *Agentic AI and autonomy: CSPs set out their strategies*, published by TM Forum last autumn.

Where and how operators are experimenting with and deploying agentic AI depends on their operational and business goals. Often the starting point is customer service. Telcos have struggled to provide affordable, satisfactory customer care for years. It involves massive transaction volumes and many repetitive tasks that are suited to automation. It has a huge, direct impact on customers' satisfaction and is expensive to provide, despite attempts to implement self-help and self-care.

Telefónica

CX AGENTS IN PRODUCTION

Telefónica's customer experience programme has six agents in production. They include a customer service contextual agent which listens to live calls and proactively sources billing records, plan details, device offers and payment guidance for human team members. The operators says this is improving the quality and speed of service interactions while reducing escalations.

QUOTE

"As the agents collectively can detect and respond to security events faster than threats can propagate, their use is likely to appeal to regulated industries, such as financial services, healthcare and government. In future, such capability could be the difference between winning and losing contracts"

Source: Philippe Ensarguet, VP of Software Engineering at Orange group

On the B2B front, a customer agent for small- and medium-sized businesses handles quotes for additional SIMs, basic billing inquiries and light troubleshooting. An order-to-catalogue agent for enterprises and partners automatically translates negotiated offers into configurations for the back-end product catalogue. An order resolution agent plus a service design agent integrate design tools with runtime network status. A root cause/correlation agent finds and fixes faults faster.

orange™

RUNNING MULTI-AGENT SECOPS

Orange is running multi-agent collaboration for complex operational scenarios. Its 5G Security Operations agents will soon enter production in a system where a log/listener agent monitors the core network for traces for security events. A knowledge-management agent is connected to knowledge bases and ticketing systems to provide context. A configuration management/signature agent interprets known security signatures and recommends responses.

These specialised agents collaborate, with each contributing domain expertise to "accelerate detection, grounding and mitigation of security issues," Philippe Ensarguet, VP of Software Engineering at Orange group, explains. The multi-agent approach avoids the risk of hallucination that can arise when single large agents try to handle all tasks.

As the agents collectively can detect and respond to security events faster than threats can propagate, their use is likely to appeal to regulated industries, such as financial services, healthcare and government. In future, such capability could be the difference between winning and losing contracts.

Orange's system exhibits "expert-like chain-of-thought and varied paths that converge to consistent outcomes," according to Ensarguet, but the operator's focus is firmly on "confidence and interpretation in non-deterministic workflows" before enabling write access to production infrastructure.

To clarify, an AI agent with read access can look at systems, for example, check network performance data, but cannot make changes – a human must act on what it finds. An agent with write access could reconfigure a network element, say, or close a ticket or update a customer account without a human approving each step first.

Orange's engineering agents for network reliability augment rather than replace expert human decision-making, providing rapid root-cause grounding and mitigation scenarios across regions.

EXPLOITING OPERATIONAL DATA

BT's strategy emphasises domain-specific – mobile, traditional voice and broadband – model training combined with leveraging cross-brand and cross-customer data in context. Séanín McCoy, Director for Data and AI in Networks at BT Business, says rather than relying on general-purpose LLMs, BT is building its own models, trained on the company's operational data, to give agents telecoms expertise.

BT's strategy is to invest in a relatively few, high-impact use cases to deepen technical capability and deliver business outcomes, rather than "scattergun experimentation".

EXPLAINER

An AI agent with read access can look at systems, for example, check network performance data, but cannot make changes – a human must act on what it finds. An agent with write access could reconfigure a network element, say, or close a ticket or update a customer account without a human approving each step first

The operator is also working through cultural and organisational factors in tandem with the technological advances. The plan is to combine deep AI science, that is, core model-building with rigorous guardrails and ethics, to deliver "AI for everyone", thereby exploiting existing platforms for productivity gains.

Its approach acknowledges that different parts of the organisation need different relationships with AI, with some building foundational capability and some leveraging tools built by others.

LEVERAGING KNOWLEDGE AND INVESTMENT

One area of focus for Vodafone is unstructured document intelligence, leveraging natural language-based search and querying to the company's knowledge sources, including PDFs, contracts, invoices and equipment documentation, for example, used by network teams and finance, with cross-functional adoption. Ahmed Hany, the operator's Big Data Chapter lead, reckons being able to find relevant information instantly from thousands of documents promises to transform operational efficiency.

Vodafone is also developing an agentic planning and optimisation system to combine the roles of data agents with documentation agents and planners' knowledge. The aim is to identify areas that potentially yield high returns on investment, assessing commercial outcomes and technical efficiency, and optimising that investment.

Deutsche Telekom

LIVE AGENTIC NETWORK MANAGEMENT

Deutsche Telekom's (DT's) RAN Guardian was rolled out on DT's live network in Germany in late 2025. It was built with Google Cloud and targets 5G's network performance because if service degrades, conventional processes for detecting, diagnosing and fixing the issue is slow, labour-intensive and typically reactive, meaning customers are affected.

RAN Guardian is a multi-agent system that monitors network parameters to predict and detect anomalies. It can either implement corrective actions autonomously, or route issues to the right team before customers experience a problem. It also scans social media and events to anticipate peaks in demand to ensure the network is ready. DT claims to be the first operator to deploy such an agent in live network management.

"With the introduction of the RAN Guardian Agent, we are the first network operator to rely on a highly developed AI agent in network management. Our development teams have thus done pioneering work in the development of AI agents for mobile networks," explains Abdu Mudesir, Board member for Product and Technology at Deutsche Telekom.

"With an intelligent interaction between our network experts and AI, we are solving specific challenges for the benefit of our customers – for the best network. And we are taking a big step towards autonomous, self-healing networks."

QUOTE

"Our development teams have done pioneering work in the development of AI agents for mobile networks. With an intelligent interaction between our network experts and AI, we are solving specific challenges for the benefit of our customers – for the best network. And we are taking a big step towards autonomous, self-healing networks"

Abdu Mudesir, Board member for Product and Technology at Deutsche Telekom

vodafone

CUTTING COSTS, BETTER OUTCOMES IN THE FIELD

Vodafone's Field Technician Assist platform was developed with implementation and integration specialist Celfocus, sponsor of this report. It engages AI across the fieldwork lifecycle.

Engineers receive contextual briefings on their way to a site. On arrival, they are guided by real-time troubleshooting recommendations drawn from historical intervention data, performance metrics and technical documentation. Remote visual support is available if needed. Closing issues is via speech-to-text input. The system continuously improves as technicians provide feedback about the recommendations.

According to this panel discussion from *The Briefing*, *AI, efficiency and profit*, the deployment has led to about 30% fewer repeated site visits, faster fault resolution and measurable improvement in customer satisfaction scores, the compound effect of which can be substantial for large telco organisations that dispatch thousands of field technicians daily

A secondary benefit is standardisation: as institutional knowledge is moved into the platform, reducing dependence on the most experienced individuals and providing consistent information. It also accelerates onboarding.

The platform is designed as a foundation for agentic AI. Future enhancements will enable autonomous decisions and proactive problem-solving in the field, without needing a technician's query to trigger a response.

In the next section we draw some conclusions and offer suggestions and recommendations for success with AI-enabled autonomous network operations (ANOps).

PARTNER CONTENT

Moving to scalable agentic autonomous networks

Partner content by CELFOCUS

Telco networks are becoming smarter, but operating them is getting harder — rising complexity, widening skill gaps, slow fault response, and untapped institutional knowledge are pushing legacy tools past their limits.

Modern network operations often behave like vast, interdependent finite-state machines — thousands of elements, countless event sequences, and subtle failure modes requiring coordinated, context-aware responses — and agentic AI closes the loop by linking high-level reasoning to real-time action.

By translating causal diagnosis into accountable remediation, orchestrating cross-domain workflows, and surfacing precise, actionable recommendations, agentic systems reduce mean-time-to-repair and human workload, encode and adapt expert playbooks to bridge skill gaps, deliver adaptive operational intelligence across sprawling topologies, enable rapid intent capture via natural language, and preserve institutional knowledge at scale.

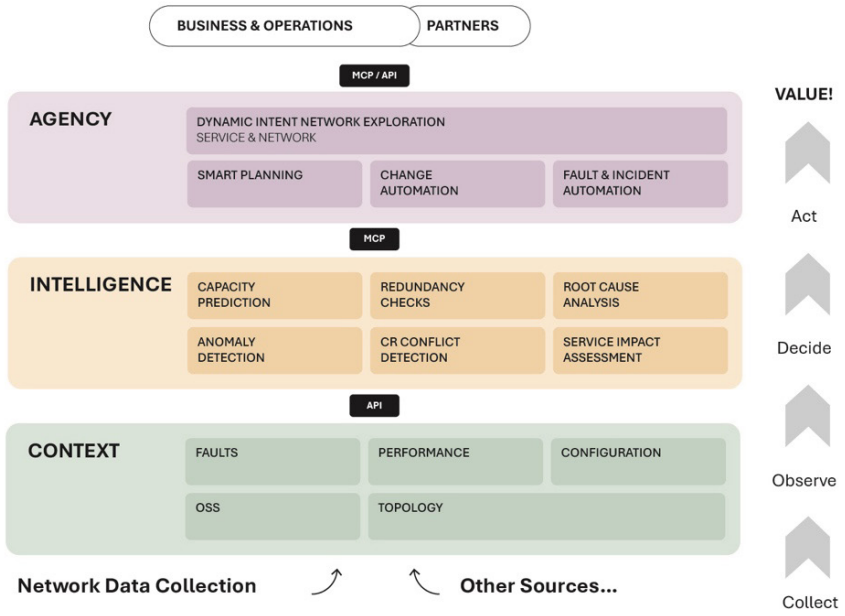

Three layers that turn raw network data into autonomous, trustworthy action

Source: Celfocus

"No agency without intelligence. No intelligence without context."

Three layers that turn raw network data into autonomous, trustworthy action"

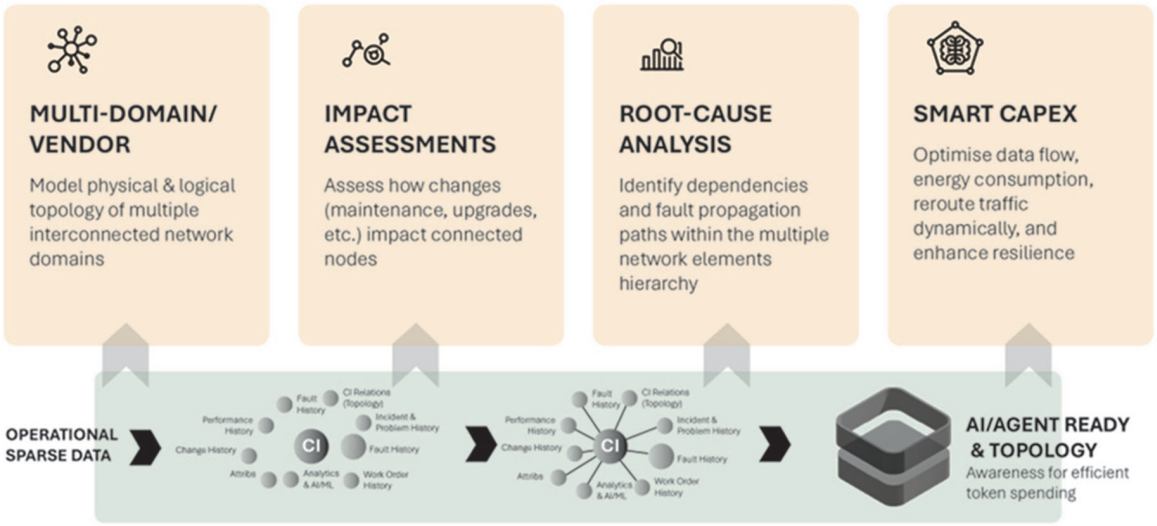

The agent's world model
Source: Celfocus

The Importance of an intelligent context-aware ecosystem

Most conversations about agentic AI in network operations jump straight to the agentic part: the prompts, the tools, the reasoning loop. That's understandable — it's genuinely exciting territory. And it matters. But it's far from the only thing that matters when you're trying to use agents to drive real, game-changing network transformation.

What gets far less attention — and arguably deserves more — is the information model underneath. The data agents are reasoning on top of. Because no matter how capable the model, an agent without a grounded, accurate view of its context isn't really reasoning. It's just expensive autocomplete, hallucinating with confidence.

That's why we advocate for a layered approach (see graphic on page 20), a substrate of structured world knowledge at the base, with intelligence (ML, LLMs, algorithms) and agency (agent frameworks) built on top of it — each layer doing what it does best.

Agents work in a loop: observe, decide, act, verify, deliver value. Every step in that loop needs a queryable, consistent, time-aware view of the network. Without that substrate, the agent rediscovers the network on every prompt, wastes effort reassembling context it should already have, and cannot reason about change. And change is 80% of what operations is.

This is where the digital twin earns its place — not as a visualisation layer or a dashboard backend, but as the agent's world model: the structured, living representation of the network the agent thinks against.

For that foundation to carry agentic workloads, it should deliver four things:

1. **Structured topological retrieval** – agents query a graph of nodes, interfaces, services and dependencies, not a pile of text. Answers are precise and repeatable.
2. **Temporal depth** – past state is a first-class query. Root-cause analysis is about change, and an agent that only sees the present is missing the most important variable in any incident.
3. **Validation surface** – changes get simulated against the twin before execution. This is what makes autonomy auditable rather than reckless, and what unlocks the high-stakes use cases: automated change management, closed-loop provisioning.
4. **Multi-vendor abstraction** – the twin normalises vendor differences once. Every agent inherits it for free, instead of each new use case re-paying the same integration tax.

Together, these four properties turn the digital twin into a foundation that makes agentic operations possible. The intelligence layer (anomaly detection, root cause analysis, impact assessment, capacity prediction) sits on top and feeds agents with decisions it can trust. The agency layer sits above that and turns those decisions into actions using multiple LLMs according to tasks at hand.

The shape of the agent layer

With a solid context layer underneath, the next question is layered above. The debate usually collapses into two unhelpful camps: pick a framework or pick a model. Both miss the point. Frameworks and models are the least durable decisions an operator will make. What matters is the shape of the agent layer itself.

- **Many small agents, not one big one.** Monolithic agents demo well and break in production. The pattern that works is a pool of

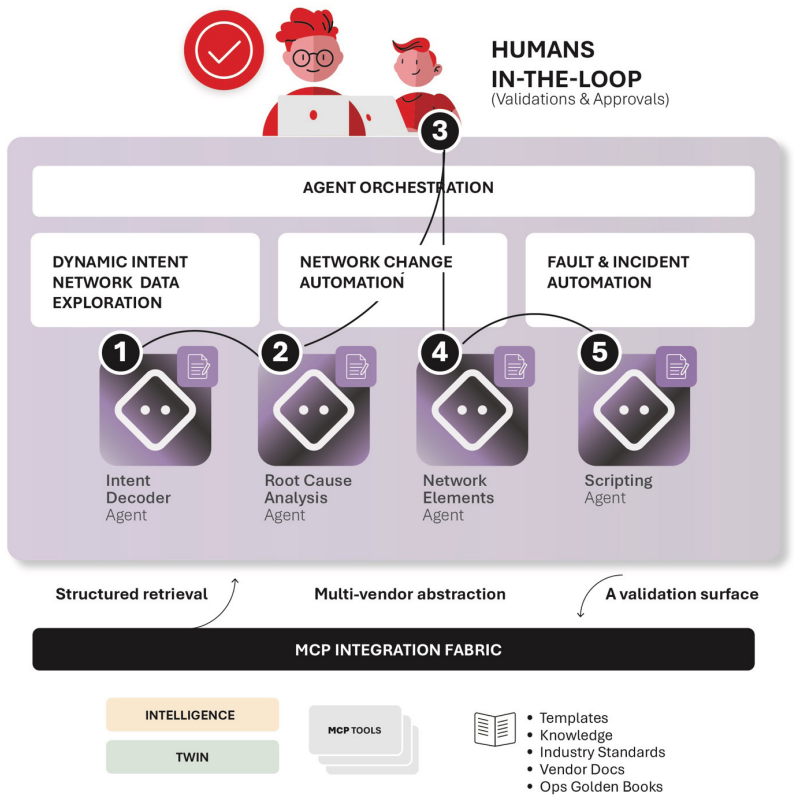

The key decision is not whether to pick a framework or a model, but the shape of the agent layer itself

Source: Celfocus

specialised agents — intent decoder, root cause analysis, network elements, scripting — each narrow enough to evaluate in isolation and replace independently. The intelligence lives in how they compose, not in any single agent.

- **Orchestration is the real work.** The orchestrator decides sequencing, state flow, failure handling and escalation. A useful test: if you can't sketch the flow on a whiteboard and explain it to an ops engineer in five minutes, it's too complex to operate.
- **Tools are the real interface to the world.** Agents act through typed functions — not directly on the network. Safety, RBAC and determinism live in the tool layer, not in prompts. Schema in, schema out; scope enforcement inside the tool; independently testable and versioned. Standards like MCP matter because they make tools portable — swap the model or framework and the tools survive. Models get cheaper. Frameworks rise and fall. A well-built tool library outlasts both.
- **Choose models for fit, not benchmarks.** Small models for routing and classification, frontier models for hard reasoning. Two rules: pin and evaluate (silent vendor upgrades are how agents quietly get worse), and stay portable (wrap model calls behind an abstraction so swapping is a config change, not a refactor).

The uncomfortable truth is that the model is rarely the bottleneck. When an agent fails in production, the cause is almost always upstream of the model — bad context, missing tools, weak orchestration. Upgrading the model just makes the failures more articulate.

The architecture, in one sentence: specialised agents, composed by an orchestrator, acting through typed tools, grounded in a digital twin, evaluated continuously and bounded by guardrails. Everything else is implementation detail that should be free to change without the architecture changing with it.

INSIGHT

The uncomfortable truth is that the model is rarely the bottleneck. When an agent fails in production, the cause is almost always upstream of the model — bad context, missing tools, weak orchestration. Upgrading the model just makes the failures more articulate.

Scaling agents without losing control

Six concerns separate agents that survive production from agents that get quietly switched off after the pilot. None are exotic. All are non-optional.

Reasoning observability – a single agent run fans out into dozens of model calls and tool invocations. "The model hallucinated" is not a root cause. Treat agent runs like distributed traces: full logging of every prompt, every tool call, every state transition, with replay so post-mortems become engineering instead of guesswork.

Data sovereignty – subscriber data, topology and CDRs cannot casually traverse a hyperscaler endpoint on another continent. Bolted on at the end, sovereignty never works. Keep data regional, obfuscate at the tool layer before anything reaches the model, and self-host the most sensitive workloads. The goal is a data-flow map that compliance can actually approve.

Cost control – agentic loops are a blank cheque written in tokens. The fix is architectural, not contractual: a well-structured twin so retrieval is precise, model routing so frontier models are reserved for hard reasoning, and outcome-level cost tracking — cost per ticket, not cost per token.

Drift management – models drift when vendors push version bumps. Twins drift when networks get reconfigured out of band. Neither shows up until something breaks loudly. Pin model versions run a golden dataset on every change, and continuously reconcile the twin against live network state.

Prompt security – the moment an agent has tools, it has a new attack surface. A malicious device banner or a crafted ticket body becomes a prompt injection vector. The first principle is the hardest: treat all tool output as untrusted data, never as instructions. Then layer on

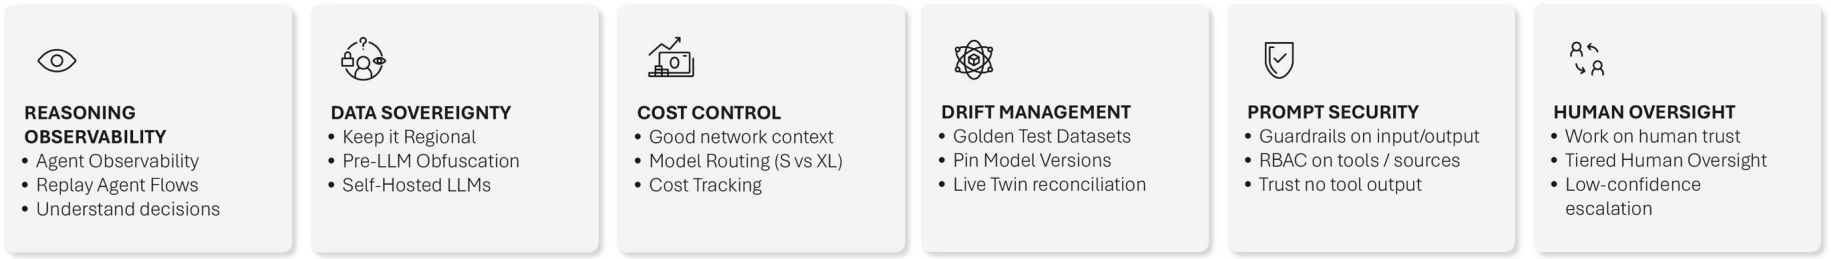

Six necessary concerns crucial to survive production

Source: Celfocus

input and output guardrails, and RBAC on tools and sources.

Human oversight – full autonomy is unsafe. Approving every step kills the value. Tiered autonomy is the equilibrium: read-only actions building human trust. Engineers will not delegate to an agent they do not understand.

The thread running through all six is the same: the model is the easy part. The hard part is the engineering scaffolding around it. The CSPs that come out ahead will not be the ones with the flashiest agents. They will be the ones who treated agentic AI as a serious engineering discipline from day one. Everything else is a prompt away.

FROM THEORY TO DEPLOYMENT: FOUR OPERATOR STORIES

The patterns above are not hypothetical. Across four Tier-1 European operators, we have been putting the same architectural discipline — grounded twin, specialised agents, typed tools — to work against very different operational problems. The use cases below span core, signaling, RAN analytics and IP transport, and together they sketch what an agentic operation starts to look like in practice.

Tier-1 UK — Multi-domain root-cause analysis in 4G/5G Core

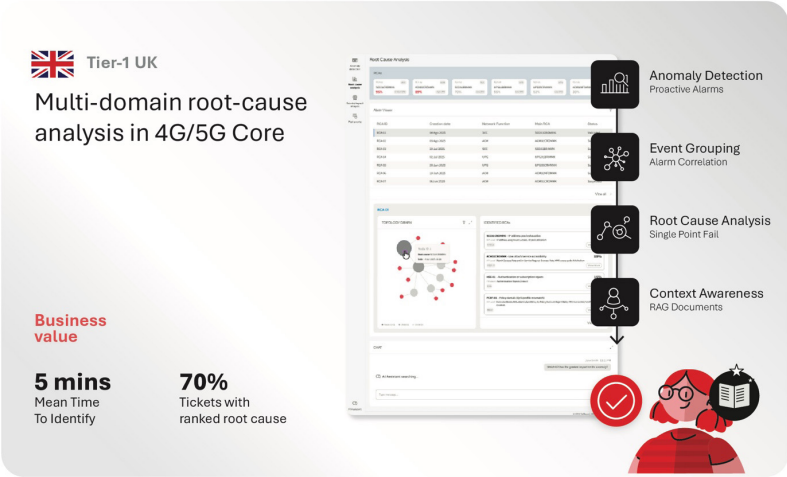

The challenge – core incidents rarely respect domain boundaries. A single ticket can span EPC and 5GC, transport and signaling, and the L2 engineer handling it ends up chasing six to ten tools to stitch a story together. Mean time to identify dominates mean time to repair — the fix is fast once you know what broke, but knowing what broke takes hours.

The agentic solution – an RCA agent pulls alarms and topology directly from the digital twin and fans out specialised sub-agents across KPIs, vendor documentation and the operator's own golden book of known failure modes. Each sub-agent contributes evidence; the orchestrator composes it into a ranked hypothesis that the engineer can accept, reject or drill into.

INSIGHT

Core incidents rarely respect domain boundaries. A single ticket can span EPC and 5GC, transport and signaling... Mean time to identify dominates mean time to repair — the fix is fast once you know what broke, but knowing what broke takes hours.

The business value – mean time to identify collapses to around five minutes, and 70% of tickets arrive at the engineer with a ranked root cause already attached. The human is still in charge — but the hours of correlation work are gone.

Tier-1 Germany — Anomaly detection in 4G/5G core signaling

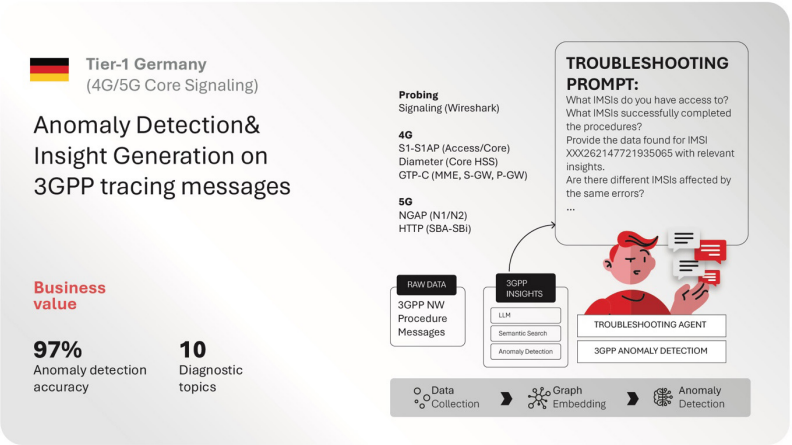

The challenge – signaling is the operator's nervous system and also one of its least tractable monitoring problems: dozens of procedures, multiple protocols, and data volumes that make manual correlation of events and messages effectively impossible.

The agentic solution – classical ML detects the anomaly; an LLM-powered agent interprets the tracing messages around it, correlates the pattern against historical incidents and turns raw signaling into a human-readable insight. The agent does not replace the anomaly

detector — it makes its output actionable.

The business value – 97% anomaly-detection accuracy and ten distinct diagnostic topics the agent can reason about end-to-end, from detection through explanation.

Tier-1 Belgium — Natural-language access to network operations data

The challenge – operational data in a modern RAN and CORE is rich, but effectively inaccessible to most of the people who need it. Extracting an insight demands deep expertise, complex tooling or a queue into a specialist team — and in operations, speed is the whole point.

The agentic solution – an engineer types something like “show me PRB utilisation for LX sites with more than 5% drop-call rate.” The agent decomposes the intent, resolves the entities against the twin, generates the underlying queries, and renders the visualisation. The specialist tool chain is still there — the agent just makes it addressable in plain language.

The business value – around 30% reduction in tooling and licensing costs as niche dashboards consolidate, and reports and dashboards that used to take hours now land in seconds.

INSIGHT

Connectivity changes on the IP backbone are exactly the kind of work that keeps network leaders awake: risky, multi-step, deeply human-reliant... every stage is a place where a small mistake becomes a large incident.

Tier-1 Spain — Closed-loop change management in IP transport

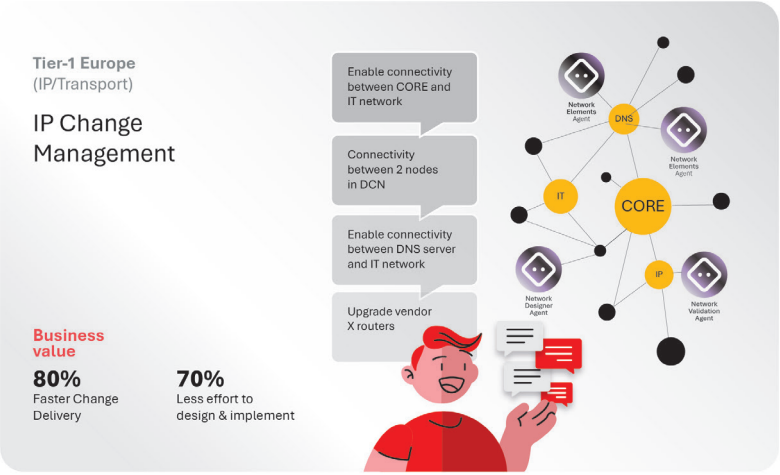

The challenge – connectivity changes on the IP backbone are exactly the kind of work that keeps network leaders awake: risky, multi-step, deeply human-reliant. Design, peer review, MOP writing, lab validation, maintenance windows, rollback planning — every stage is a place where a small mistake becomes a large incident.

The agentic solution – the engineer gives the agent an intent — for example, "extend this L3VPN to PE-Y." The agent designs the change against the digital twin, generates the scripting, simulates the result, executes inside the maintenance window, and either commits or rolls back based on validation. This is the use case the validation surface in the twin was built for: autonomy that is auditable because it is grounded.

The business value – 80% faster change delivery and 70% less effort to design and implement — with a rollback story that compliance and the NOC can both sign off on.

Taken together, these four deployments make a concrete point. The architectural choices in the first half of this article are not abstractions. They are what lets the same underlying platform handle an RCA in London, a signaling anomaly in Frankfurt, a natural-language query in Brussels and an IP change in Madrid — without four different agent stacks, and without four different integration bills.

Conclusion: the hard part is the foundation

Agentic AI is having its moment in telecoms, and a lot of that attention is landing in the wrong place. The interesting question for CSPs is not which model to pick, which framework is fashionable, or how many agents a vendor can fit into a demo. The interesting question is whether the ground underneath the agents is solid enough for them to stand on.

A high-fidelity, time-aware, multi-vendor digital twin. A specialised-agent architecture with typed tools and a thought-through orchestrator. Observability, cost control, data sovereignty, drift management, prompt-injection defences and tiered human oversight — all treated as first-class engineering concerns, not afterthoughts. None of this is glamorous. All of it is what separates a pilot that excites the board from an operation that quietly runs the network at three in the morning.

INSIGHT

[Operators] are closing the loop deliberately: starting with read-only insight, graduating to reversible action, and promoting an agent to irreversible change once it has earned the trust of the engineers... That progression is slower than the hype cycle would like, and it is exactly the right speed

The operators we are working with are not chasing autonomy for its own sake. They are closing the loop deliberately: starting with read-only insight, graduating to reversible action, and promoting an agent to irreversible change only once it has earned the trust of the engineers who work alongside it. That progression is slower than the hype cycle would like, and it is exactly the right speed.

The prize at the end is real. Incidents resolved before customers notice. Change delivered in hours instead of weeks. Institutional knowledge that survives the next retirement wave. Engineers freed from correlation drudgery to do the work only they can do. But none of it arrives by accident, and none of it arrives by prompt engineering.

It arrives by treating agentic AI as what it is: a serious engineering discipline, built on a serious foundation. The CSPs that internalise that now will spend the next few years turning the hard part into an advantage. Everything else really is a prompt away.

CONCLUSION

10 Recommendations for strategic progress

The journey to AI-enabled autonomous network operations (ANOps) has begun, but operators must be prepared for the long haul. Successful deployment and real return on investment will require disciplined execution, grounded in operational and economic reality.

1

Don't expect a big short-term impact on P&L

Telefónica's Caroline García warns that short-term impact on P&L is limited. A more realistic goal is holding operational costs steady as the complexity of networks accelerates with 5G SA, network slicing and other new types of service. To control costs, Telefónica applies FinOps discipline to AI, testing different model architectures including LLMs, SLMs, fine-tuning and RAG so it can choose the most cost-efficient approach for each use case, rather than defaulting to the most 'capable' one.

As mentioned on page 5, Charlotte Patrick calculates that two-thirds of the value of autonomy within the network has already been extracted. The harder to realise cases that remain might well be uneconomic as sophisticated agentic systems require substantial investment in data platforms, integration, knowledge management and ongoing AI operations. Also, payback periods are potentially measured in years rather than months. Cost and benefit need to be weighed carefully.

QUOTE

"[Starting in] areas where they use some level of automation and AI already, for example, chatbots in customer care, billing, marketing....will accelerate time-to-value while reducing risk"

Marina Koytcheva, Research Director at STL Partners

2

Start narrow, build confidence, scale strategically

The most successful approaches focus on specific, well-defined problems with a clear ROI. Marina Koytcheva, Research Director at STL Partners recommends that operators start in "areas where they use some level of automation and AI already, for example chatbots in customer care, billing, marketing." Building on existing foundations, rather than starting from scratch, will accelerate time-to-value while reducing risk.

Projects must be selected carefully. Patrick cautions against the risk of technical debt associated with embarking on "lots of little projects with agents that get low hanging fruit" only to realise later "that actually we needed to do something different really." Tactical shortcuts now could require costly rewrites later.

3

Data quality is foundational

In a recent interview Scott Petty, Group CTO at Vodafone, says one of the toughest things about adopting GenAI then agentic AI has been sticking to the common IT and data architectures that are used across the group, resisting hype-driven pressure to incorporate various vendors' AI solutions in the interests of speed and short-term expediency. Note that 'shared ops' were highlighted as a critical success factor in Vodafone's growth strategy in the group's earnings presentation for in November 2025 for the first half of the current financial year.

Swisscom's Düsener sees data quality as a prerequisite for "massive" investment in AI with agentic AI trials "imminent".

Verizon in the US is building through an AI maturity curve, from predictive to prescriptive, generative, then agentic. High-quality, multi-modal data is the prerequisite for each step, according to Kalyani Sekar, Chief Data Officer, Verizon.

Analyst Patrick stresses that data hygiene must be "immaculate", so that all downstream consumers receive accurate, context-aware inputs. This means standardising data formats and taxonomies across disparate systems; building federation capabilities for cross-domain access; implementing continuous quality scoring; and establishing governance with clear accountability.

Regarding the creation and maintenance of knowledge graphs, ontologies and curated documentation, the real challenge is organisational. Who owns curation? How are updates validated? What happens when documentation conflicts with operational reality? These questions must be resolved before knowledge foundations are scaled.

4

Keeping humans-in-the-loop while building trust

Human oversight is essential during early deployments and for validation. Vodafone's deployment of Field Technician Assist (see page 18) started with a controlled pilot and offline data. Then it progressed to a limited live field test with selected technicians, providing iterative feedback on AI recommendations. From there was extended to a full production roll-out.

José Palma, Head of Network Strategy at the Portuguese operator MEO, recommends a similar path to that it took for self-organising network (SON) optimisation in its RAN. First the platform runs in an open loop; human operators review and validate every recommended configuration change before it is applied to the network. Once humans verify the results, the change is switched to a closed loop where the system ingests telemetry, determines the level of optimisation and reconfigures the network autonomously. RAN optimisation that once took days now takes minutes.

Orange's approach – agent proposes, human approves, agent executes, human monitors, human involvement gradually reduces as confidence grows – is a sound model.

INSIGHT

Orange's approach is that human involvement gradually reduces as confidence grows – a sound model

agent proposes

human approves

agent executes

human monitors

5

Assess and measure added value customers will pay for

At Swisscom, Mark Düsener's approach is to explore what value AI enables that customers will pay more for. Choosing such services has been the operator's highly successful guiding principle for a decade: he points out that 80% of Swisscom's revenues come from services that didn't exist 10 years ago.

Proximus Ada, a Belgian centre of excellence, has a measurement framework that combines AI and cybersecurity which provides a useful model for understanding the value delivered via AI, according to Emmanuel de Hermicourt de Grunne. He is Global Head of Digital and Zero Trust Operations at Proximus Ada.

The framework is based on customer metrics like the improvement in customers' satisfaction from searching FAQs, more right first time service calls and employees' productivity. The framework also monitors tools like GitHub Copilot and economic metrics such as the avoidance of capex and opex, and revenue generation and protection.

Ahmed Hany, Vodafone's Big Data Chapter Lead, says measurement

frameworks must be defined before deployment, as without clear metrics, it is hard to separate success from failure, which impedes learning and improvement.

He says AI deployment can be tied to direct financial outcomes using agentic AI in smart capex allocation to analyse network performance and commercial data across regions. Agentic systems can identify where investment will generate measurable returns within a defined payback period. "We can measure the usage, whether it's traffic usage, or whether we are selling significantly more products," Hany says.

Where direct revenue baselines are harder to establish, Vodafone applies utilisation rates and customer satisfaction scores as proxies for adoption and value.

6

Clouding the ANOps issue

As agents move from simple tasks to complex orchestration, complexity increases exponentially. BT Business' Colin Bannon identifies concerns around supply chains.

He thinks there is risk in the concentration of demand in AI hardware, but also insufficient resilience within hyperscalers' single zones and he stresses the importance of platform-level resilience. According to Bannon, agent systems that are production-grade require a level of infrastructure resilience beyond current cloud architectures.

QUOTE

"[Agentic systems can identify where investment will generate measurable returns within a defined payback period.] We can measure the usage, whether it's traffic usage, or whether we are selling significantly more products"

Ahmed Hany, Vodafone's Big Data Chapter Lead

7

Modelling reliability and non-determinism

GenAI's tendency towards hallucination remains problematic in operational environments. Operators are aware of this: for example, Orange ensures its models exhibit "expert-like chain-of-thought and varied paths that converge to consistent outcomes," explains Philippe Ensarguet, VP of Software Engineering at Orange.

This is a powerful antidote, but requires new approaches to testing and validation. Traditional test plans that assume deterministic outputs don't work for systems that reason differently each time while reaching similar conclusions.

8

Security, governance and trust

As Ensarguet states, security and sovereignty are "non-negotiable" in telecoms. Challenges include the Model Context Protocol (MCP) which needs integration that's enterprise-grade secure, as well as authentication, authorisation and accountability across agents. The handling of sensitive data and compliance requirements are onerous too.

BT Business' "sovereign-by-design" approach, as outlined by Bannon, includes "models trained and platforms operated within the UK" with "local platforms with UK-run compute, data residency, and freedom from external legal interference," such as the US' CLOUD Act.

The Agentic AI Summit at MWC 2026 in Barcelona included discussions of new dangers as AI gains greater access privileges. Emerging attack vectors including memory poisoning, which is corrupting an agent's context to manipulate behaviour and compromising privilege so that agents gain unauthorised access through theft of credentials or escalating privileges.

9

Organisational and cultural barriers

People-based considerations need to be front and central to AI ambitions. Natali Delić, Chief Strategy and Digital Officer at Telekom Srbija, notes, "AI projects are not IT projects." As agents proliferate, business units will increasingly work alongside what she calls digital employees, that is, AI agents embedded in operational teams. The ramifications of these hybrid teams need to be thought through.

For Delić, a necessary shift from task-oriented processes to domain-level AI governance, with business ownership at its centre, is what will set AI-native organisations apart from those who simply run AI tools.

Patrick believes operators must "develop a rolling programme to build trust and manage employee worries around their job security." Workers, she says, fear displacement while managers are worried about losing control and technical staff have questions about AI's reliability. These human factors could constrain progress more than technical limitations if not addressed.

QUOTES

"AI projects are not IT projects"

Natali Delić, Chief Strategy and Digital Officer at Telekom Srbija

"The real fun [hard work] starts when you look at services or cross-domain scenarios"

Afnan Ahmed, Director of Technology Strategy and Architecture at Telenor

10

Setting realistic expectations for all stakeholders

There needs to be a grounding of expectations for ANOps. It is not a magic wand. There is a long way to go before any telco can demonstrate self-healing and self-optimisation in significant parts of their network, Patrick warns. They need new technologies and painstaking work must be done to develop knowledge and data layers, then there's the complexity of integrating closed loops across processes, and difficult to answer questions about return on investment.

At Telefónica, despite its ambitions and considerable advances with nine high-value, Level 4 use cases live across Spain, Brazil and Germany, the average autonomy level was 3.05 at the company's mid-2025 assessment. Telefónica has publicly stated that it is targeting an average of 3.75 by 2028 and Level 4 across its core domains by 2030. Universal autonomy across all operations remains years away and is perhaps unachievable as technology and business goals change all the time.

Afnan Ahmed, Director of Technology Strategy and Architecture at Telenor, is involved in TM Forum's Autonomous Networks (AN) project, focusing on the challenges and requirements for reaching Level 4 autonomy. He notes that most claims around Level 4 achievements are not domain-wide, but more likely to be within a specific domain, such as energy efficiency in the RAN for 4G and 5G. "The real fun [hard work] starts when you look at services or cross-domain scenarios," he adds. A Level 4 organisation remains a long way off.

Successful operators will be those that recognise AI-enabled network autonomy as a capability to be built systematically over years, not a product that can be bought and deployed in months.

ME MOBILE EUROPE | INDEPENDENT RESEARCH