



www.pipelinepub.com

Volume 22, Issue 11

AI has a visibility problem: Why Trust Will Decide the Future of Autonomous Telecoms

By: [Guy Lupo](#)

Not long ago, telecom operators were asking whether they should adopt artificial intelligence (AI). Today, the question is very different.

AI is already here. It is embedded in network operations, customer service, security functions, software development and business processes. New AI agents, models and tools appear seemingly every week, while employees and suppliers continue to experiment with new capabilities. The challenge facing operators is no longer adoption. It is visibility.



A useful analogy comes from an unlikely source: Pokémon.

The current AI landscape often feels like a world filled with Pokémon. They appear everywhere. New ones emerge constantly. Different teams are using them for different purposes. Some are approved and governed. Others may not be. Leaders are left with the hard part: finding them all, understanding what each one can do, and building a complete, living inventory.

For telecom operators, that question has significant consequences. Every AI agent, model, or application potentially introduces new decisions, new data dependencies, and new operational risks. Before organizations can truly benefit from AI at scale, they need to know what AI exists within their business, how it behaves, and whether it can be trusted.

This is becoming one of the defining challenges of the industry's journey toward autonomous operations.

The evolution of AI leadership

The rise of AI is also reshaping leadership within telecom organizations. A few years ago, Chief Data Officers often found themselves responsible for managing data platforms while struggling to demonstrate business value. AI initiatives, meanwhile, frequently sat within isolated AI programs, IT

teams, or specific business functions. Those boundaries are rapidly disappearing.

Today's AI systems depend upon access to data, applications, network capabilities, and security controls. An AI agent cannot deliver meaningful outcomes if it operates as a disconnected tool sitting outside the operational environment. It needs to interact with enterprise systems, understand data, and execute actions across the organization.

This convergence is creating a new generation of leadership roles. Chief AI and Data Officers are increasingly being positioned closer to executive decision-making and are assuming broader responsibilities than technology enablement alone. Their remit is expanding to include governance, risk management, security, sovereignty, and operational accountability.

In some organizations, this evolution is already producing a new title: Chief AI and Trust Officer. As AI becomes more influential in operational decision-making, leaders will inevitably be asked the same question by boards, regulators and executives: Can we trust it to run our business? Answering that question requires more than technical expertise. It requires evidence, governance and accountability.

Automation vs. autonomy

The telecom industry has spent decades pursuing automation. Automated workflows, policy-based controls, and intelligent operations have become familiar parts of modern network management. Autonomy represents something different. Automation focuses on executing predefined actions. Autonomy requires systems to make decisions within established boundaries and adapt to changing circumstances while remaining aligned with organizational objectives.

Many organizations describe advanced AI as automation on a larger scale. In practice, true autonomy requires a fundamental shift in how decisions are delegated to machines. The industry is moving from models that rely on humans reviewing every action – “human in the loop” – toward environments in which humans establish objectives, policies, and controls while autonomous systems determine how those outcomes are achieved – “human in charge of the loop.” That transition cannot occur without trust.

Organizations must trust the data used to train AI systems, the models generating recommendations, and the behavior of agents interacting with critical infrastructure. They must also trust that they can intervene, audit, and explain decisions when required. Without that trust, AI remains confined to pilots and demonstrations. With trust, it becomes possible to scale autonomy across complex operational environments.

The data foundation

Many of the most successful AI deployments in telecom share a common characteristic: trusted data foundations. When organizations discuss AI performance, attention often focuses on model selection, agent frameworks, and emerging technologies. Yet the quality of outcomes frequently depends on something less glamorous but far more important: the underlying data. Reliable data creates predictable outcomes. Structured data creates repeatable outcomes. Governed data creates auditable outcomes. In practice, organizations that have successfully introduced greater autonomy into fault management and assurance consistently find trusted data to be the primary enabler of success.

Conversely, poorly governed data introduces uncertainty. AI systems become harder to explain, harder to validate, and harder to trust. As organizations move toward autonomous operations, data governance is therefore no longer an administrative exercise. It becomes a strategic capability. The path to trusted AI begins long before model deployment. It begins with confidence in the information

that powers decision-making.

The security imperative

Trust cannot be discussed without addressing security. Threat actors are already benefiting from AI. Cybercriminals are using automation and AI to accelerate reconnaissance, identify vulnerabilities, and increase the speed of attacks. The industry has long argued for a shift from human-speed to machine-speed security operations, and AI is increasingly making that possible.

But operationalizing AI-driven security in telecoms is not straightforward. Operators manage some of the world's most critical digital infrastructure, where changes are carefully controlled, and security architectures carry significant governance requirements. The challenge is not whether AI can improve security outcomes. It is how to integrate it safely without creating new vulnerabilities or undermining existing protections.

Technology may be advancing rapidly. Operational change rarely moves at the same pace.

The governance challenge of agentic AI

The emergence of agentic AI introduces another layer of complexity. Unlike traditional software applications, agents may make decisions, interact with systems, and collaborate with other agents to achieve objectives. As organizations deploy increasing numbers of these autonomous capabilities, governance becomes significantly more difficult. This is the same Pokémon challenge at enterprise scale: before organizations can govern agentic AI, they must first find, identify, and catalog every moving part.

The questions are fundamental: how many agents exist across the organization, what models are they using, what data can they access, which decisions are they authorized to make, and how are they being monitored? These are governance questions, but they are also business questions. The answers influence security, compliance, operational resilience, and customer trust. They also influence costs. As AI usage grows, organizations need greater transparency into how models are consumed, where workloads are processed, and how resources are allocated.

The industry therefore faces a new requirement: managing AI ecosystems rather than individual AI solutions. Success will depend less on deploying the latest model and more on establishing effective mechanisms for governance, orchestration and oversight.

Closing the gap

The scale of the challenge is measurable. Published industry research surveying 130 operators worldwide found that [72 percent](#) of communication service providers believe their AI is trustworthy, but only 14 percent can produce evidence to prove it. Governance frameworks exist on paper. Operational proof does not. Closing that gap requires progress on three fronts.

Organizations need a standardized approach to sourcing, operating, and governing AI models. Whether models come from hyperscalers, vendors, or in-house teams, a common control layer is what makes model behavior consistent, auditable, and lifecycle-managed.

Data must be treated as a governed product: discoverable, structured and subject to clear contracts for ownership and use, so that AI agents can access it safely and consistently without manual intervention at every step. Agentic systems also need explicit security frameworks: common policy languages and guardrails that define authorized behavior, enforce boundaries, and generate

machine-readable audit trails at scale.

Regulation is hardening this agenda into a requirement. The EU AI Act requires high-risk AI systems to meet audit and governance standards, and proposed European legislation would introduce enforceable sovereignty assurance levels for public sector AI procurement. McKinsey's 2026 AI Trust Maturity Survey reinforces the picture: organizations with explicit accountability for AI trust achieve consistently higher maturity scores, and investment in responsible AI is strongly associated with realized business value. Trust, the research concludes, is increasingly a business enabler rather than a compliance exercise.

Across the industry, operators and their partners are already building the collaborative, standards-based foundations this requires. The direction is clear. But individual effort will not close the gap at the speed or scale the industry needs.

Trust as an operational capability

Historically, trust has often been treated as a soft concept. Organizations make claims about being trustworthy. Vendors describe trusted platforms. Leaders communicate their commitment to responsible AI. Those discussions remain important, but they are no longer sufficient. In an autonomous environment, trust must become measurable.

Organizations need to demonstrate why an AI system made a particular decision, with full visibility into the data and models involved. They need evidence that controls function as intended and confidence that governance mechanisms hold as systems evolve. Trust therefore becomes an operational capability rather than a marketing message.

This mirrors previous advances within telecom: reliability, security and network performance all became measurable disciplines, not aspirational qualities. AI trust is following a similar trajectory. The future will belong not to organizations that simply claim their AI is trustworthy, but to those that can demonstrate it.

The road to autonomous telecoms

The telecommunications industry is moving steadily toward greater autonomy. AI agents will become more capable, models more specialized, and autonomous systems will assume increasing responsibility across networks, IT environments, and customer operations. The complexity of managing these environments will grow with them.

This creates an important paradox. The more AI organizations deploy, the more important governance becomes. The more decisions machines make, the more confidence leaders require in the systems making them. Autonomy and trust are not separate objectives. They are inseparable.

The industry often frames autonomy as a technology challenge. The real challenge is simpler and more fundamental. Can organizations understand what AI they have? Can they govern it effectively? Can they explain its behavior? Can they trust it to operate within acceptable boundaries? The operators that answer those questions successfully will be best positioned to unlock the full value of autonomous operations.

The race to autonomous telecoms is therefore not fundamentally an AI race. It is a trust race. And it starts with the basics: knowing what AI exists across the business, how it is being used, what data it relies on, and whether its decisions can be explained, governed, and trusted. Without that foundation, autonomy cannot scale safely.