



www.pipelinepub.com

Volume 22, Issue 11

Unlocking AI's Full Potential with Humans in the Loop

By: [Miguel Carames](#)

AI is already a critical component of today's telecom networks. However, utilization does not imply that operators are ready to hand over the control and operation of their networks to AI agents. According to the [TM Forum](#), 46 percent of CSPs surveyed said they have deployed AI in production for specific functions within IT & digital systems and networks, with slightly fewer in operations and service management (41 percent) and security (39 percent). We are at the tipping point of entering the AI-enabled era.



AI has the potential to reduce costs, drive growth, differentiate customer experience, and ensure secure, reliable operations. In fact, it's imperative to achieve the promise of autonomous networks and deliver a "Zero X" experience – zero wait, zero touch, zero trouble through AI-enabled self-management, configuration, optimization, and security. According to the [World Economic Forum](#), reaching Zero X operations is estimated to be worth \$794 million per year in cost savings, efficiency gains, and revenue generation opportunities for CSPs.

But we are not there yet. AI still has a way to go before it can prove that it is trustworthy to be handed the keys to the operations. AI-related operational and security incidents remain a common challenge for CSPs around the globe. In fact, the TM Forum found that **only 4 percent of CSPs reported no AI-related incidents in the past 12 months**. The report revealed that the most AI-related risk came from:

- 47 percent said that AI generated inaccurate outputs that affected operations or customers in the past year
- 28 percent had experienced incidents of prompt injection, whereby hackers disguised malicious inputs as legitimate prompts
- 16 percent had suffered AI model tampering or "poisoning."

These statistics underscore why a human-in-the-loop approach – where people review, validate, and

guide AI decisions - isn't optional: it keeps operations efficient and reliable while catching risks such as configuration drift (where configurations gradually slip from their intended baseline) and model hallucinations (inaccurate or nonsensical output), problems that surface only through sustained human oversight and deep domain knowledge. Human-in-the-loop isn't just about risk mitigation, though; it is also necessary to advance the journey toward autonomous operations. By providing context, validating AI recommendations, and continuously refining the models and the reference data that grounds those models, human experts help AI become more accurate, reliable, and trustworthy over time.

Here are the top five areas where human-in-the-loop is critical to unlocking AI's full potential.

1. Risk Reasoning

Where the value of human-in-the-loop is required is where analysts validate the AI's reasoning, correct its blind spots, and address concerns around bias, discrimination, liability, and compliance. As a result, they gain true real-time visibility into their risk posture so that they can respond to incidents faster and with more confidence and can develop a unified risk strategy across their operations, from billing and revenue collection to fraud management, revenue assurance, and compliance.

2. Testing and Monitoring Root Cause Analysis

That same interplay between AI and human judgment shows up just as clearly on the network side, where data volumes are exploding. According to [PWC](#), data consumption is expected to increase from 3.4 million petabytes (PB) in 2022 to 9.7 million petabytes (PB) by 2027. With the help of AI, it is possible to turn the sheer volume, variety, and velocity of raw data into instant KPI reporting and benchmarking, flagging network issues, and triggering alarms before customers ever notice. However, true Root Cause Analysis (RCA) automation is what transforms testing and monitoring from reactive troubleshooting into proactive assurance. Unleashing the potential of RCA automation still requires humans in the loop to validate AI-flagged anomalies, apply operational context to complex, multi-vendor issues, refine the balance between deterministic and probabilistic AI that ensures that even the more nuanced scenarios are properly analyzed and the best possible recommendations described. The opportunity for CSPs is to optimize network performance — ensuring superior Quality of Experience (QoE) and keeping operators ahead of industry trends.

3. Revenue Assurance Evidence Analysis

The same principle carries over into revenue assurance. AI is changing revenue assurance from a reactive, data-combing exercise into a conversational and interactive one. An analyst can now ask "why did roaming leakage spike?" and get back a contextual explanation with supporting evidence and recommended next steps, as AI routes the query to the right dataset or tool and interprets the results. A revenue assurance agent, for instance, can explain why revenue leakage increased following a network configuration change, tie the spike to common industry failure scenarios, flag the controls affected, and point analysts to where they need to focus their investigation.

Humans are still needed to validate that an AI's explanation holds up against the actual transaction and network data, catch and investigate cases where a leakage pattern doesn't cleanly match a well-known scenario, and apply judgment about business context, customer relationships, or regulatory exposure that the model has not been trained to weigh yet. The result is AI functioning not as an isolated tool, but as an intelligent advisor that leans on human oversight to ensure the evidence matches the conclusions before any decisions are made.

4. Democratizing Insights for Data Monetization

Human oversight isn't only about catching errors, though — it's also about controlling who gets to ask the questions in the first place. As network usage outpaces ARPU growth, the spotlight firmly focuses on the widening monetization gap. For years, the industry's ambitions for data monetization have ridden on hopes of AI and advanced analytics. Years later, AI still isn't delivering — and part of the reason is simpler than it looks. The tools were created for engineers and were built to perform specific technical tasks. But it's the business users, not engineers, who are asking the monetization questions.

The advancements of GenAI and AI agents mean that AI can be put in the hands of business users across the organization. Where it once was the domain of IT and programmers, AI can now be queried in plain language by the marketing manager pursuing a churn insight, the product lead pricing a new bundle, or the finance analyst hunting down a revenue leak - without waiting weeks for a data science team to build a bespoke model or dashboard. This is what democratization actually means: not fewer engineers, but far more people empowered to ask monetization questions and get answers directly from the network and customer data that exists under the hood. That shift matters because the monetization gap was never really a data problem or even an AI problem — it was an access problem. Telecom operators have been sitting on rich troves of usage, roaming, and behavioral data for years, but the insight stayed locked behind technical interfaces that only a small team of experts could operate and understand, so the people closest to the revenue decisions were the last to see the answers.

Democratizing AI closes that gap by putting the analysis in the hands of the business users who own the monetization questions, turning AI from a specialized engineering tool into an everyday decision-making resource — and giving operators a real opportunity to close the gap between network usage and ARPU. This, too, is human in the loop — only here, the human is the business owner asking the question, not a data scientist standing between them and the answer.

5. Expanding Autonomous Operations

Bring risk reasoning, root cause analysis, evidence analysis, and democratized insight together, and the natural next step comes into view: autonomous operations. According to WEF, 65 percent of working hours within a CSP can be transformed by large language models (LLMs) - where 36 percent have a high potential for automation, and 30 percent for augmentation. That scale of change is exactly why human oversight can't be an afterthought: it's the mechanism that makes autonomy trustworthy, not a temporary stopgap. Every reviewed or corrected AI decision feeds a playbook the AI draws on for future autonomous action — a playbook built on frameworks that map risks, controls, and evidence to standard requirements for consistent governance.

So, why is human-in-the-loop important for autonomous operations? Automating the network isn't enough on its own — autonomous operations also need intelligence that understands the business, capable of telling a routine technical glitch apart from a genuine commercial threat, explaining the risk it poses, and coordinating the right response. Human insight is what's still needed to judge what actually matters to the business, spot where exposure is building, and adjust accordingly.

The stakes are real: AI failure can mean prolonged outages, BSS/OSS breakdowns that halt revenue collection, or compromised data incidents that trigger regulatory fines and reputational damage. So, while many are putting the foot on AI's accelerator, CSPs still need to keep their hands on the wheel as they build autonomy incrementally, with humans validating each step, so AI earns greater responsibility over time.

Ultimately, it all comes back to one simple truth. No two days in a telecom network look the same,

and that unpredictability is precisely why AI needs the human in the loop. Humans supply the skill layer that raw network data can't: the nuance, the context, the judgment calls that turn a signal into a decision. Achieving autonomous operations tomorrow depends on human input today, and for that input to count, AI must be usable by the people who actually understand the business, not just the engineers who built it. In practice, this means rules and machine learning detect the risk, agentic AI accelerates the understanding and execution, and humans hold onto control and governance throughout. And the payoff compounds. Every interaction and correction sharpens the system further, so explanations, recommendations, and prioritization keep getting more accurate, more attuned to what your organization needs, and ultimately more automated.

Not for distribution or reproduction.