

SEPTEMBER 2026

Bringing AI to the Data: Trusted Autonomous Networks for Telecom with Cloudera and NVIDIA

Stephen Catanzano, Principal Analyst

Abstract: Telecom operators face flat growth, declining average revenue per user (ARPU), and saturated markets. The path to sustainable margins runs through autonomous networks, the rare lever that simultaneously reduces service delivery costs and improves customer experience. But the data that makes autonomous networks possible is massive, distributed, and sits on-premises. It's also not practical to move the data to where AI is traditionally run, many times due to cost, data sensitivity, and governance. The answer is to bring AI to the data. And no operator will hand its most valuable asset to an AI system it cannot audit. Cloudera, accelerated by NVIDIA, delivers governed, explainable AI to data wherever it lives, enabling telcos to achieve the performance and trust they need to automate the network itself.

The automation imperative

The telecom industry has pursued automation for decades, but the economy now makes it an existential imperative. Growth has stalled. The World Economic Forum projects roughly 0.3% CAGR, with flat-to-declining ARPU through 2027,¹ even as markets saturate and customer acquisition costs climb. Operators are shifting their primary growth levers from acquisition to retention, churn reduction, and customer lifetime value. Sustainability depends on doing two things at once: delivering an exceptional customer experience and drastically reducing service delivery costs. Autonomous networks are a rare capability that pulls both levers simultaneously, which is why they have moved from a long-term aspiration to a near-term imperative.

Where network AI breaks down

The obstacle of autonomous networks is architectural. Network data is generated and stored where the network lives, on-premises and at the edge, while AI has traditionally been run in centralized environments far from that data. For the telecom industry, closing that distance by moving data is unrealistic.

73% of organizations are actively implementing or piloting sovereign AI solutions. This emphasizes organizations' need to bring AI to their data.

¹ Source: World Economic Forum White Paper, [The Strategic Role of Telecom Providers Across the AI Value Chain](#), March 2026.

Network data is both massive and highly distributed. According to Cloudera, a single North American operator processes roughly 6 petabytes of network data every day, about 22 billion network events, spread across hundreds of sources spanning the radio access network (RAN); the core, transport, IoT, and connected devices; and business support and customer systems. Moving that volume of data to a central location is financially and technically impractical, and the physical distance between centralized AI and the network introduces latency that undermines real-time operations. For closed-loop autonomous networks, this challenge is compounded by the bi-directional flow of information: data must flow to the AI system for inference, and AI recommendations based on that inference must then be sent back to the network for action. This round-trip communication not only introduces latency that degrades real-time inference and response capabilities but also adds significant reliability and operational risks to mission-critical network functions.

There is a second, quieter barrier: trust. Telcos will not hand the keys to the network over to a black-box model they cannot explain, which is precisely why customer experience use cases have outpaced network use cases to date. Autonomous networks require a different approach: Bring AI to the data and make its decisions accountable. In recent Omdia research, respondents were asked what best describes their organization's use of, or interest in, sovereign AI. Overwhelmingly, 73% responded that their organizations are actively implementing or piloting private AI solutions/self-hosted AI solutions² to not move their data to AI and so they can control how AI works with their data.

From assisted to autonomous

The operators making progress are deploying a layered approach rather than a single model. Traditional machine learning provides a real-time foundation for detecting anomalies such as distributed denial-of-service (DDoS) attacks, preventing fraud, and enabling predictive maintenance that catches faults before they cause outages.

87% of organizations agreed that AI agents' capabilities will enable them to solve problems and address use cases that were previously impossible.³

Generative AI adds synthesis and explanation, powering customer support and turning raw signals into narratives an engineer can act on. Agentic AI orchestrates these capabilities into closed-loop workflows that detect, diagnose, and act.

The trajectory is a deliberate progression: from humans in the loop to on the loop to out of the loop (optionally, depending on the use case), each step ceding more autonomy as trust is earned. That trust is not a matter of model accuracy alone. It depends on governance and explainability (i.e., the ability to see why an AI system has reached a decision before granting it control over the network). Autonomy and accountability advance together—or not at all. Governance practices are also changing rapidly. Omdia research found that nearly seven in ten organizations (69%) described the governance challenge posed by agentic AI as very or extremely significant. Only 4% said their governance is already designed for real-time autonomous processes. That

² Source: Omdia Research Report, [The Great Repatriation: Why Enterprises Must Act Now to Ensure AI Sovereignty](#), July 2026.

³ Source: Omdia Research Report, [AI Agents: The Game-changing Generative AI Use Case](#), August 2025.

means 96% of organizations are operating agents against a governance model that was not built for them.⁴

Use case: Closed-loop network automation

Consider a failed cell tower. Today, resolution is largely manual. Alarms fire, engineers triage across systems, and mean time to repair stretches while service degrades. In autonomous networks, ML detects the anomaly the moment performance drifts, generative AI assembles the context (recent events, dependencies, and probable cause) into a clear diagnostic summary, and an agent retrieves the relevant data, evaluates remediation options, and either executes the fix or recommends one, with a human on or out of the loop, depending on the operator's risk tolerance.

Because inference runs in the network, the loop closes in real time rather than waiting on a round trip to centralized infrastructure. The outcomes are tangible: faster mean time to repair, lower operational costs, and an improved customer experience because problems are resolved before subscribers notice them.

Use case: Autonomous energy optimization

Energy is one of the highest and most controllable costs in the network, with the RAN accounting for a substantial share of power consumption. Static, manually tuned energy policies leave savings on the table because operators fear degrading service. Autonomous optimization changes the equation. Using the NVIDIA Blueprint for intent-driven rAN energy efficiency on the Cloudera platform, operators can continuously and automatically adjust the network to minimize energy use while holding service-level commitments intact, capturing cost and sustainability gains that manual approaches cannot safely reach. The use of digital twins to verify decisions made by the agentic AI system is an essential component of such autonomous systems. The same closed-loop foundation extends to use cases such as dynamic 5G network slicing, where network capacity is allocated in real time to match demand.

The platform: Cloudera and NVIDIA


Solving the gap between where data lives and where AI runs requires bringing AI to the data rather than moving the data to the AI. That is the core of the combined Cloudera and NVIDIA platform.

Cloudera provides the trusted data foundation. Its open data lakehouse unifies distributed sources into a consistent, real-time view, while streaming ingestion and stream processing—with MiNiFi filtering noise at the edge—deliver live network health and performance. Hybrid deployment enables operators to run each stage of the AI lifecycle where it makes sense: training in the cloud and inference in the network. Cloudera's Shared Data Experience delivers consistent security, governance, and end-to-end lineage from ingestion to consumption, with the model explainability that regulated operators require to trust AI with the network. Cloudera's depth here is well established, with 9 of the world's top 10 telcos running on the platform.

⁴ Source: Omdia Research Report, [Mastering Governed Autonomy: From Policy Blind Spots to Competitive Edge](#), July 2026.

As shown in **Figure 1**, NVIDIA supplies the acceleration that makes on-premises and edge AI performant: NVIDIA CUDA-X library, cuDF for GPU-accelerated data processing, NVIDIA NeMo Retriever for enterprise retrieval-augmented generation in AI Workbench, NVIDIA NIM microservices for scalable inferencing, and NVIDIA Nemotron as the model foundation for Agent Studio's agentic workflows, augmented by NVIDIA Blueprint for telco network configuration. Together, the platform keeps sensitive data in its jurisdiction, enabling the private and sovereign AI that telcos increasingly view as a strategic opportunity.

Figure 1. Cloudera AI: Build and run private AI



Source: Cloudera and Omdia

Conclusion

The telecom industry's margin pressure is not going to ease, and autonomous networks are the most powerful response available, one that cuts the cost of service delivery and lifts customer experience at once. The barrier has never really been ambition. It has been an approach that keeps AI far from the network data it must act on and a justified refusal to trust decisions that cannot be explained. Cloudera, accelerated by NVIDIA, resolves both: It brings governed, explainable AI to the data, wherever that data lives, with the performance real-time networks demand. That combination is what finally enables operators to extend automation beyond the call center and into the network itself. Omdia recommends that telecom leaders engage with Cloudera to explore how this platform can accelerate their path to autonomous, trusted network operations.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the “Omdia Materials”) are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together “Informa TechTarget”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

Get in touch: www.omdia.com askananalyst@omdia.com

