



www.pipelinepub.com

Volume 22, Issue 8

Enabling Multi-Vendor Autonomous Assurance with Agentic AI

By: [Kenan Jarah](#)

Challenges in Driving Autonomy in Telecom Networks

Telecommunication networks are undergoing a fundamental transformation driven by cloud-native architectures, virtualization, distributed infrastructure, and increasingly complex multi-vendor ecosystems. Modern environments comprise highly distributed network functions, microservices, and radio elements that continuously generate large volumes of telemetry across multiple operational domains.



Ensuring reliable service delivery requires Communications Service Providers (CSPs) to continuously monitor and correlate performance metrics, fault events, and service indicators across domains, including the Radio Access Network (RAN), the core network, and cloud infrastructure. While the availability of data has significantly increased, the ability to consistently interpret and act on that data across domains remains a challenge.

Despite advances in analytics and automation, operational workflows remain constrained by fragmented toolchains, domain-specific silos, and reliance on manual investigation and escalation processes. Root-cause identification across domains remains slow, visibility across multi-vendor environments is often limited, and automation is typically confined within domain boundaries. These structural constraints directly affect operational efficiency and service reliability.

At the same time, the rapid advancement of Machine Learning and Generative AI introduces new capabilities. Machine Learning enables deterministic analysis of large-scale telemetry, including anomaly detection and forecasting, while Agentic AI introduces the ability to coordinate actions, reason across systems, and enable goal-driven execution.

Rethinking Operational Architecture: Toward Distributed Intelligence

A key observation emerging from modern telecom environments is that operational complexity is not driven solely by data volume, but by the distribution of intelligence across domains and systems. Traditional approaches have often relied on centralized data aggregation to enable analysis; however, in distributed and multi-vendor environments, such approaches introduce challenges related to scalability, governance, and latency.

The architectural approach explored in this article shifts the focus from centralizing data to coordinating intelligence. The architecture minimizes reliance on full data centralization by combining a structured data foundation with localized processing across systems. This allows intelligence to remain within each domain or vendor environment, while coordination is achieved through structured interaction.

This model enables each system to contribute its domain-specific capabilities within its own context, while higher-level orchestration combines these contributions into coherent operational outcomes. The result is a more scalable and adaptable approach to managing distributed telecom environments.

Way forward with Multi-Agent Operations

In response to these challenges, an Agentic AI-based operational framework can be designed to enable coordinated network assurance across domains and vendors.

At the core of this approach is the ability to correlate and analyze events across multiple domains, including RAN, core, and cloud infrastructure, without relying on isolated tools. This supports a unified interpretation of performance, fault, and service data, improving the ability to identify and resolve issues efficiently.

Multi-vendor collaboration is treated as a fundamental requirement. The architecture enables interoperability between independent systems while preserving their autonomy. Vendor-specific intelligence remains within each system but can be accessed through structured interaction, reducing reliance on manual escalation and improving resolution cycles. The framework also emphasizes distributed intelligence. By minimizing data movement and enabling localized execution, it supports scalability while addressing governance constraints. At the same time, it introduces an intent-driven operational model, where workflows are dynamically determined based on objectives rather than predefined sequences.

Another key aspect is the integration of deterministic and agentic intelligence. Deterministic models provide statistically grounded insights, while agentic components coordinate investigation, escalation, and execution. This separation ensures both analytical accuracy and operational flexibility.

Bringing the Framework to Life

The implementation of this framework is structured across multiple layers, including Codap as the data foundation and GenAie as the agentic orchestration layer.

Codap: Data Foundation and Semantic Layer

Codap serves as the foundational data layer, providing a structured abstraction over heterogeneous telecom data sources. It ingests and organizes data from multiple operational systems, including performance management, fault management, topology and inventory systems, and customer experience data.

Through normalization and enrichment, Codap ensures consistency across domains and vendors. Data is organized using a medallion architecture, progressing from raw ingestion to AI-ready insights. This layered approach supports traceability, reproducibility, and controlled data transformation.

Codap also introduces a data product model, where curated datasets are exposed for consumption by analytical models and agents. These datasets are enriched through an ontology layer that defines relationships between network entities, enabling contextual interpretation and cross-domain correlation.

In addition, Codap integrates a knowledge base containing domain expertise, operational rules, KPI definitions, and known failure patterns. This enhances the ability of analytical and agentic components to interpret data and align insights with operational scenarios.

While Codap provides a centralized abstraction for structured and governed data, the broader architecture minimizes reliance on full data centralization by enabling localized processing across systems.

GenAie: Agentic AI and Orchestration Layer

GenAie represents the agentic layer of the architecture, responsible for orchestrating workflows, coordinating agents, and enabling intent-driven decision-making across domains and systems.

At the center of GenAie is the Supervised Agent, which acts as an orchestration entity capable of responding to multiple types of triggers, including analytical signals, user-driven interactions, and external system requests. Based on these inputs, the Supervised Agent constructs workflows, selects relevant agents, and manages execution across systems.

Unlike traditional rule-based orchestration, this model is adaptive. Workflows are not predefined but dynamically constructed based on context, allowing the system to respond effectively to complex, multi-domain scenarios.

GenAie includes domain-specific agents such as Performance Management, Fault Management, and Trouble Ticketing agents. Each agent performs specialized tasks, including KPI correlation, alarm validation, and workflow integration. These agents operate within a task-oriented execution model, where workflows are decomposed and executed through coordinated interactions.

A key capability of GenAie is Agent-to-Agent (A2A) collaboration. Through A2A communication, the system interacts with external vendor agents, enabling domain-specific analysis to be performed within vendor environments. Only structured insights and results are exchanged, allowing each system to maintain control over its data while contributing to coordinated workflows.

Protocols Enabling Coordination

The effectiveness of this architecture depends on structured interaction mechanisms that enable coordination across systems.

The Model Context Protocol (MCP) provides a standardized way for agents to access data and tools in a context-aware manner. Instead of interacting directly with raw data sources, agents query structured data products and invoke analytical capabilities through a controlled interface.

Complementing MCP, Agent-to-Agent communication enables collaboration across system boundaries. In this model, interactions are defined by intent rather than procedural instructions. One agent specifies the objective, and the receiving system determines how to fulfill it within its own domain.

This dual interaction model supports both internal coordination and external collaboration, enabling workflows to span multiple systems while minimizing reliance on full data centralization.

Governance and Controlled Autonomy

As automation increases, governance becomes essential. The architecture incorporates policy-driven execution, ensuring that all actions are performed within defined constraints. Different execution modes are supported, ranging from fully autonomous operation to human-in-the-loop scenarios where approval is required.

Traceability is built into the system, with all interactions, decisions, and actions recorded. This ensures that workflows can be audited, reconstructed, and validated. Decision-making is also explainable, with supporting evidence and reasoning preserved.

This approach enables controlled autonomy, balancing automation with operational oversight.

Summary and Industry Direction

The architecture described in this article represents a shift in how telecom operations can be structured. By combining deterministic analytics, agentic orchestration, and structured inter-agent communication, it enables distributed intelligence to be coordinated across domains and vendors.

Rather than relying on full data centralization, the approach combines a structured data foundation with distributed execution, allowing intelligence to remain within each system while enabling collaboration through intent-based interaction. This addresses challenges such as data silos and fragmented workflows without requiring extensive integration.

The evolution of such architectures will depend on continued industry alignment, particularly in the standardization of Agent-to-Agent communication and interoperability frameworks. Initiatives from organizations such as TM Forum and GSMA will play a key role in enabling consistent adoption across vendors and platforms.

Ultimately, the transition toward AI-native telecom operations depends not only on more advanced analytical models but on the ability to coordinate intelligence across systems, domains, and vendors in a controlled and scalable manner. The combination of multi-agent systems, structured data foundations, and standardized interaction mechanisms provides a practical path toward this objective.