TechnologyAdvice

INDUSTRY TREND REPORT

Agentic AI in Action:
Building Autonomous Networks for the 5G and 6G Era

SPONSORED BY

UST

CONTENTS

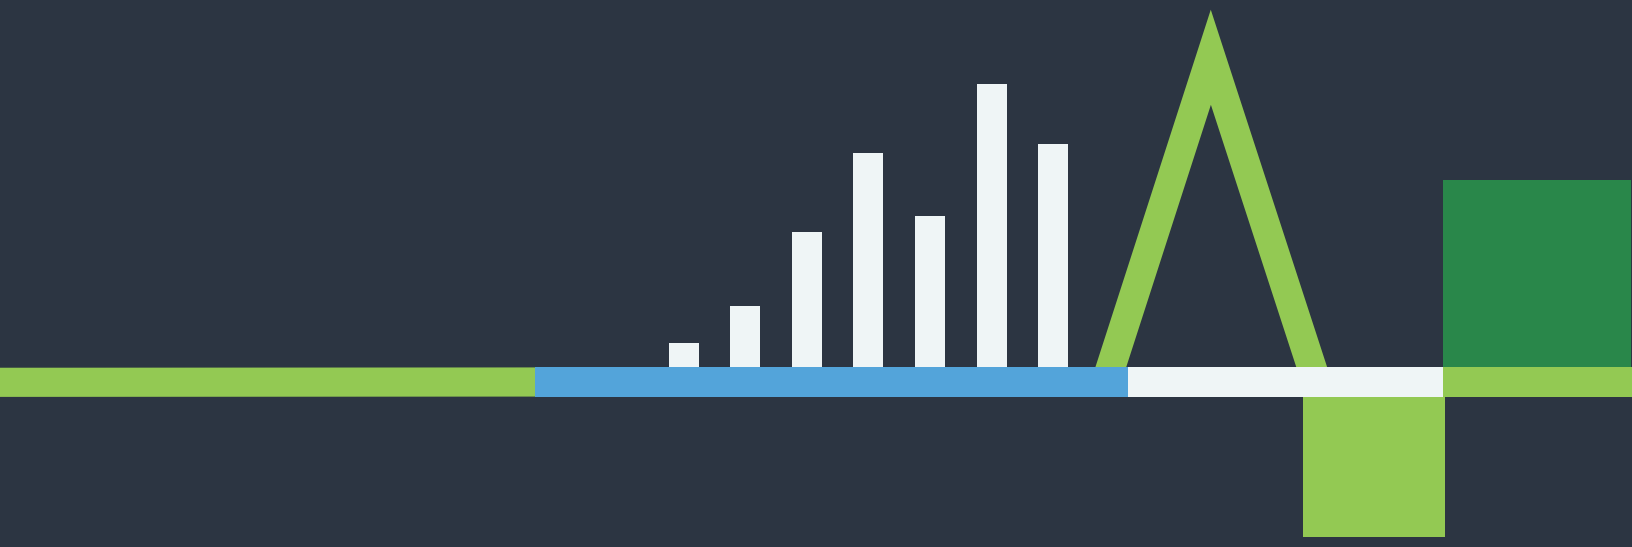

Agentic AI in Action:
Building Autonomous Networks for the 5G and 6G Era

> The most exciting part of telecom's transformation today is that we're finally seeing the long-held vision of "networks that run themselves" take shape. Agentic AI — systems that can orchestrate, perceive, decide, and act autonomously — is becoming the intelligence layer that connects automation with true autonomy.

RAMKUMAR VENKATESALU
DIRECTOR OF TELCO ENGINEERING AT UST

EXECUTIVE SNAPSHOT

Agentic AI: The Shift from Automation to Autonomy (2023-2025)

Agentic orchestration is now operational in Tier 1 telecom environments. AI agents are evolving beyond static automation to interpret intent, coordinate across domains, and act with real-time autonomy. Human-on-the-loop oversight is replacing manual triage, enabling self-directing networks.

What Matters to Operators

Operators are prioritizing initiatives that not only improve operational efficiency but also drive measurable impact on customer experience and business outcomes.

Opex efficiency: >99% reduction in alert noise and 30–40% fewer manual tasks (SmartOps), enabling leaner operations and faster resolution times.

Energy optimization: 71% of operators have reduced network energy usage via AI-enabled dynamic resource control, delivering both cost savings and sustainability benefits.

Service velocity: Custom slices and private 5G networks can now be deployed in hours, not weeks, accelerating enterprise time-to-value and creating new revenue streams.

Revenue and experience protection: Proactive fault prediction and SLA assurance workflows help reduce churn by improving uptime and maintaining service quality.

YOUR 90-DAY PATH TO FIRST VALUE

WEEKS

1-2

Identify one high-noise or high-cost use case (e.g., fault triage).

3-5

Align data inputs and define guardrails (what AI can/can't act on).

6-8

Launch bounded pilot with measurable KPIs (precision, MTTR, energy).

9-12

Align data inputs and define guardrails (what AI can/can't act on).

What Must Be True First

Before agentic AI can deliver real autonomy, operators must address the foundational challenges of today's telecom environments, characterized by fragmented tools, siloed domains, and overlapping vendor systems. Successful implementation depends on aligning data, policy, and control across these heterogeneous environments.

To enable effective AI decision-making across fragmented environments, the following capabilities must be in place:

Unified data fabric: Normalized, real-time data across OSS/BSS, CNFs, and cloud infrastructure.

Policy + intent models: Standards-driven frameworks (e.g., ETSI ZSM, TMF) to align automation with business goals.

Risk controls: Guardrails for rollback, explainability, and human-on-the-loop oversight.

Cross-domain visibility: End-to-end telemetry integrated for closed-loop control.

Key Executive Decisions

1. **Start where value is proven:** Alarm triage, prediction, energy management.
2. **Define your governance model:** AI explainability, approval workflows, kill-switches.
3. **Assess data readiness:** Can you normalize telemetry across domains?
4. **Map your operating model shift:** From manual ops →AI-augmented NOC teams.
5. **Choose pilot partners:** Look for proven deployments with measurable impact.
6. **Evaluate organizational readiness:** Identify required skill sets (e.g., AI operations, model oversight) and plan for upskilling or role evolution.
7. **Decide on build vs. partner vs. extend strategy:** Weigh internal development, external platforms, or enhancements to existing OEM tooling based on speed, flexibility, and long-term integration needs.

99.3%

alert noise reduction and easing NOC workload through intelligent triage1

PROVEN VALUE

Agentic AI is already delivering measurable operational improvements in live 5G environments, especially in assurance, predictive operations, and energy management. Early deployments show that operators can reduce noise, accelerate resolution, and shift from reactive workflows to proactive control without requiring full autonomy from day one.

In one Tier 1 environment, UST SmartOps condensed more than 100,000 daily alarms into a small, correlated set of actionable tickets, achieving 99.3% alert noise reduction and easing NOC workload through intelligent triage[1]. In another deployment, prediction agents delivered 95% precision and forecasted approximately 75% of alarms at least 24 hours ahead of occurrence[2]—providing teams more time to prevent service degradation. Importantly, these predictions were delivered under a human-on-the-loop governance model, where AI agents flagged likely incidents with supporting context, while operations engineers retained final control over response actions.

Industry benchmarks reinforce these outcomes. Capgemini reports that early autonomous network programs have delivered up to 20% improvement in operational efficiency and 18% reductions in network operating costs, as operators automate orchestration, assurance, and optimization workflows[4]. Energy management is also emerging as a high-return use case: Capgemini found that 71% of operators have reduced network energy consumption over the past two years through AI-supported automation and dynamic resource management[4].

UST's SmartGenie multi-agent system demonstrates how these capabilities work together in practice. It combines agents that forecast potential issues, support engineer workflows, and coordinate field and resource actions, while a supervising meta-agent maintains alignment and coordination across the network[3].

USE-CASE PORTFOLIO

USE CASE	VALUE FOCUS	TIME TO VALUE	DATA READINESS	RISK	OWNER
ALARM CORRELATION & TRIAGE	OPEX	4-8 WEEKS	HIGH	LOW	NOC LEAD
ALARM PREDICTION (PROACTIVE OPS)	OPEX / RESILIENCE	6-10 WEEKS	MODERATE	LOW	NOC LEAD
ENERGY OPTIMIZATION	COST SAVINGS	6-8 WEEKS	HIGH	MEDIUM	OPS LEAD
SLA/QOS ASSURANCE (PRE-DEGRAD.)	REVENUE RETENTION	8-12 WEEKS	MODERATE	MEDIUM	SERVICE OPS
DIGITAL TWIN-BASED TESTING	RISK REDUCTION	8-12 WEEKS	MODERATE	MEDIUM	ENGINEERING / QA

These use cases demonstrate where agentic AI is already delivering business value today—especially when paired with human-on-the-loop guardrails and measurable KPIs.

GOVERNANCE AND RISK

Achieving autonomy in live network environments requires strict oversight and control mechanisms. UST deployments and industry frameworks support governance models that ensure transparency, safety, and explainability across automated decisions.

Human-on-the-loop operations: AI agents provide recommendations or initiate changes, but human operators approve actions until confidence thresholds are met.

Explainability and traceability: Every AI-initiated action is auditable, and intent-policy alignment is documented across the lifecycle.

SmartOps enforcement layer: UST SmartOps includes built-in guardrails, approval workflows, and policy enforcement capabilities, ensuring AI agents operate within defined boundaries and regulatory expectations.

Digital twin validation: Proposed network changes are tested in simulation environments before production deployment, reducing operational risk.

Security posture: Agentic systems are hardened via zero-trust architecture, anomaly detection, and real-time isolation of risky behaviors.

These controls provide the accountability, oversight, and safety mechanisms needed to move from automation to autonomy in high-stakes telecom environments.

FOUNDATIONS: DATA FABRIC, INTENT, AND POLICY

Normalized data fabric

Autonomous AI systems require complete, real-time visibility across OSS/BSS, CNFs, and cloud infrastructure. A unified data fabric aggregates telemetry into a normalized model for reasoning and decision-making.

UST's Telco360 platform operationalizes this foundation by creating a unified data fabric across the network. It integrates SNMP, NETCONF, OSS/BSS APIs, Kubernetes, and other sources, giving AI agents full visibility across the radio, transport, core, and cloud layers. This enables situational awareness, reasoning, and autonomous decision-making at scale — capabilities essential to the transition from automation to autonomy[5].

Intent-policy integration

Business-level goals must be translated into policy and executable actions through models like ETSI ZSM and TM Forum's Intent-Based Management framework. These standards guide how AI interprets objectives and maintains compliance.

[1] UST. SmartOps Alert Noise Reduction Case Study. 2024. (Internal reference)
[2] UST. SmartOps Alarm Prediction Case Study. 2024. (Internal reference)
[3] UST. SmartGenie Multi-Agent Framework Overview. 2024. (Internal reference)
[4] Capgemini Research Institute. Networks with Intelligence: Why and How the Telecom Sector Should Accelerate Its Autonomous Networks Journeys. 2024.
[5] UST. Telco360 Architecture Diagram and Capabilities. 2024. (Internal reference)
[6] Hexa-X. Hexa-X Architecture for B5G/6G Networks — Final Release (D1.4). 2023.

Decision-making capability

Agentic systems perceive conditions, reason across domains and constraints, and act autonomously. This "perceive → reason → act" model is what enables real-time coordination — but it depends on clean data, clear policy, and orchestrator readiness.

6G IMPLICATIONS: WHAT TO DO NOW VS. WATCH

The shift toward fully autonomous, AI-driven networks is already underway. But autonomy at scale won't be achieved by technology alone — it will require operators to address critical integration, trust, organizational, and security challenges along the way. Agentic AI serves as a bridge: delivering measurable value in today's 5G networks while laying the operational and cultural foundations for 6G.

Integration first: Why autonomy requires clean, connected data

Agentic AI systems rely on accurate, real-time data across domains — but many networks remain siloed across OSS, BSS, CNFs, and legacy platforms. This fragmentation limits AI's ability to reason and act.

DO NOW:
INVEST IN DATA NORMALIZATION AND TELEMETRY INTEGRATION

WATCH:
CONVERGENCE OF AI TRAINING AND REAL-TIME ORCHESTRATION PIPELINES IN FUTURE 6G DESIGNS

The UST Telco360 platform addresses these integration barriers by aggregating data from SNMP, NETCONF, REST, OSS/BSS APIs, Kubernetes, cloud systems, and other EMS/NMS sources. The result is a normalized data fabric that enables consistent training, inference, and real-time action across domains[5].

Building trust through governance and simulation

Even with measurable accuracy, operators must trust how agentic systems reach conclusions — especially in high-stakes domains like assurance or traffic management.

DO NOW:
INTRODUCE DIGITAL TWINS AND EXPLAINABILITY TOOLS IN CRITICAL DOMAINS

WATCH:
POLICY MANDATES FOR AUDITABILITY AND SAFETY IN AI-LED ORCHESTRATION

Explainable AI (XAI) and network digital twins help establish this trust. XAI enables engineers to understand the rationale behind AI actions. Digital twins allow teams to simulate and validate decisions in virtual environments before they reach live networks. These tools are already in use in early 6G trials, where simulations help confirm whether a decision is safe and effective.

DO NOW:
BUILD INTERNAL CAPACITY AROUND AI GOVERNANCE AND NETWORK SIMULATION

WATCH:
INDUSTRY-WIDE SHIFT TO EXPERIENCE-LEVEL AGREEMENTS AND AI-LED ASSURANCE

The culture shift: From direct control to AI-augmented operations

As autonomy expands, network operations will increasingly rely on data science, model validation, and intelligent oversight. Engineers must evolve from direct control to managing agentic systems that operate independently but transparently.

This transition isn't only technical — it's cultural. Successful operators will retrain teams to collaborate with intelligent agents, manage trust thresholds, and maintain human-on-the-loop safeguards.

DO NOW:
ADOPT ZERO-TRUST PRINCIPLES AND CONTINUOUSLY VALIDATE AI AGENTS

WATCH:
REGULATORY FRAMEWORKS REQUIRING AI EXPLAINABILITY AND THREAT ISOLATION

Security, risk, and regulation in agentic systems

As networks become more software-defined and open, they also become more exposed. AI agents can detect and isolate anomalies faster than humans — but only if they're secure by design.

Zero-trust architectures and AI-powered anomaly detection are emerging as must-haves. These treat every agent, device, and interaction as untrusted until verified, supporting resilience as networks scale.

The 6G opportunity: Intelligence by design

Unlike 5G, where intelligence was added incrementally, 6G is being designed as an AI-native system. Programs like Hexa-X envision networks where digital twins, AI-driven orchestration, and real-time optimization are embedded into core infrastructure[6].

UST anticipates that agentic AI will serve as the bridge to this future. Operators deploying agents today — in alarm triage, SLA assurance, and energy optimization — are gaining the maturity, data readiness, and operational experience they'll need to lead in 6G.

DO NOW:
SCALE AGENTIC AI IN TARGETED 5G DOMAINS

WATCH:
STANDARDIZATION OF DIGITAL TWIN OPS AND FULL AI-NATIVE ORCHESTRATION IN 6G

Strategic outlook

The challenges to full autonomy are real — but solvable. They require coordination across technology, policy, and workforce strategy. Operators that act today to normalize data, build trust through validation, and adopt agentic tools in key workflows will not only improve current performance, but also establish a foundation for next-generation, self-directed networks.

Agentic AI is no longer a concept, it is a working model that helps operators scale 5G with greater precision, flexibility, and resilience. And it's the most pragmatic stepping stone toward the intelligence-first architecture 6G will demand.