AvidThink

AI for Networking: Charting the Agentic Waters

Why Context, Governance, and Trust Define the Next Wave of Autonomy

RESEARCH BRIEF



Table of Contents

Executive Summary 1
Introduction: From Pipe Dreams to Deluge to Open Waters 2
The Wider Current: What Coding Agents and Desktop Copilots Taught Networking 3
What Changed in 2026: Our Scorecard 9
The 2026 Toolkit: Predictive, Generative, Agentic, and the Harness 11
Context: Data, Source of Truth, Ontology, and Digital Twins 16
Governance: Policy, Guardrails, Identity, and Audit 19
Trust: Verification, Explainability, and Rollback 20
Interoperability: MCP, A2A, A2A-T, and the Open Agentic Ecosystem 22
A Cross-Domain Perspective: Service Provider, Enterprise, and Data Center 23
Barriers, the Human Element, and Recommendations 24
Appendix: Vendor Profiles A-1

Research Briefs are independent content created by analysts working for AvidThink LLC. These reports are made possible through the sponsorship of our commercial supporters. Sponsors do not have any editorial control over the report content, and the views represented herein are solely those of AvidThink LLC. For more information about report sponsorship, please reach out to us at **research@avidthink.com**.

About AvidThink

AvidThink is a research and analysis firm focused on providing cutting-edge insights into the latest in infrastructure technologies. AvidThink's coverage includes cloud and AI infrastructure, 5G and mobile network infrastructure, enterprise and data center networks, private wireless, edge computing, SD-WAN, SASE, ZTNA, and infrastructure security. Our clients include Fortune 500 enterprises, hyperscalers, tier-1 communications service providers, as well as innovative startups. AvidThink's research has been quoted by Forbes, the Wall Street Journal, Light Reading, Fierce Networks, Mobile World Live, and other major publications. Visit AvidThink at **www.avidthink.com**.

©2026 AvidThink LLC. All Rights Reserved.

RESEARCH BRIEF

AI for Networking: Charting the Agentic Waters

Why Context, Governance, and Trust Define the Next Wave of Autonomy

Executive Summary

In 2026, agentic AI performs real-world tasks in network operations, but in a far more constrained way than the 2025 announcements implied. Today, agents investigate, correlate, and recommend across enterprise, carrier, and data center networks. In a handful of operator-built carrier systems, they can close the loop unattended, but only within well-bounded scenarios. Reaching this stage didn't depend on a better foundation model, but on improved harnesses. The same control patterns that made coding agents useful have moved into networking: curated context available to the agent, action governance, a deterministic execution layer under the agent, and built-in token budgets and cost metering.

This year's edition focuses on context, governance, and trust. Context, including the source of truth, short/long-term memory, ontology, and digital twin, now determines an agent's accuracy. End users must decide who builds and owns that context layer. Governance, including agent identity, registries and gateways, policy, and audit, determines whether and how an agent can act.

Meanwhile, until sufficient trust is established, every agentic system will be limited to pre-authorized actions with a human in the loop. And strong governance and good agent performance are key to earning trust. The good news is that we have constrained tasks where agents take autonomous action, and closed-loop systems are delivering ROI, especially in carrier networks.

The industry is approximately where we said it would be: two years into a five-to-seven-year journey to full autonomy. The barriers of 2024 remain, from lack of trust and messy data to missing skills and unclear ROI. Thus, our advice in 2026 includes investing in data clean-up and building a strong context layer. To improve trust, ensure you have a deterministic execution layer beneath the agents for reliability. Govern agents like users, and build that governance early, before encountering agent sprawl. Employ digital twins as both a source of context and a verification method. To control costs, budget tokens the way you budget headcount. Grade your autonomy levels scenario by scenario, measure KPIs before and after, and keep a human in the loop on anything that has an expensive undo function. The question for the next twelve months is not whether agents can run networks. It is whether organizations can build the context, governance, and trust to let them.

©2026 AvidThink LLC. All Rights Reserved.

Introduction: From Pipe Dreams to Deluge to Open Waters

Looking Back to 2024 and 2025

Our 2024 report, "Pipe Dreams and AI Realities," caught the beginning of the generative AI (GenAI) wave for networking operations and automation. At that time, predictive AI had already proven itself for multiple tasks, including anomaly detection, capacity planning, and energy optimization across enterprise and carrier networks. Early GenAI provided a useful interface to comb through reams of incoming data. GenAI, coupled with retrieval-augmented generation (RAG), enabled the first generation of context-based interactions by pulling in semantically relevant sections of knowledge bases and operations manuals, letting users "chat" with context and perform early troubleshooting and basic configuration.

Our 2025 report, "Harnessing the AI Deluge," captured the next wave of reasoning models and, with them, a flood of agentic AI framework announcements from major networking vendors and numerous startups. In that report, we outlined an early, general multi-layered architecture on a unified data foundation, orchestrated through the Model Context Protocol (MCP) and the Agent2Agent protocol (A2A). We also identified trust and agent sprawl as potential barriers, and predicted a 12–18-month race to prove ROI.

What has become clear in 2026 is the importance of giving models the right context, and wrapping agents in governance harnesses, before letting them act autonomously.

What has become clear in 2026 is the importance of giving models the right context, and wrapping agents in governance harnesses, before letting them act autonomously. The agent incidents that dominated the fall 2026 news cycle show that ensuring alignment and adding guardrails to agent harnesses is challenging. This slows the path to full autonomy, which requires organizations and human operators to trust agents enough to let them close the loop in network operations.

This year's report therefore focuses on the main themes of context, governance, and trust, while also updating you on the current state of the market. We changed the preposition from "in" to "for" in the report title to make it clear we are covering AI's operational impact on networks.

What Our 2026 Report Covers

Based on our research at conferences covering multiple domains (telecommunications operators, enterprises, and data centers), as well as briefings and ongoing conversations with both end users and vendors, we've put together updated content for the 2026 edition of the AI for Networking report:

- What coding agents have taught networking
- Looking back: a brief scorecard of our 2024 and 2025 predictions
- The 2026 AIOps (AI for IT operations) toolkit and the evolving architecture
- Context, governance, and trust as the keys to autonomy
- Latest in interoperability (MCP, A2A, A2A-T) and integration
- Comparing advances across service provider, enterprise, and data center domains
- Barriers to adoption and our recommendations
- Update on select AI network operations and automation vendors

©2026 AvidThink LLC. All Rights Reserved.

Who Should Read What

Much of this year's evidence is carrier-led, as carriers seem eager to publicize gains on their journey to autonomous networks; enterprises and data centers are more private about their results. Nevertheless, we believe in cross-pollination opportunities between domains (even though they face different types and scales of challenges), so understanding the progress across domains can be instructive. If you're in a hurry, here's what we suggest you focus on:

- **Enterprise network teams**: the value is likely in the platforms you own today. Agents, context, and guardrails arrive bundled with incumbent networking vendor platforms: Cisco Cloud Control, Extreme Platform ONE, HPE Mist and Aruba Central subscriptions, Arista CloudVision, and similar offerings from other network vendors. The hard work is in brownfield integration with IT service management (ITSM) and source-of-truth systems, setting autonomy parameters with a small team, and understanding how AI usage is metered. Start with the Toolkit, including its token economics subsection, and read the All Domains and Enterprises recommendations.
- **Service providers**: if you're a tier-1 or even a tier-2, you will likely build or own the orchestration layer and the context plane, or integrate with your existing operations and business support systems (OSS/BSS) vendor. Your challenge is to scale observability and governance so agents outnumber people on staff. Start with the Context and Governance sections, and read the All Domains and Carriers recommendations.
- **Data center and cloud teams**: many of you are ahead of the enterprises, relying on GitOps, continuous integration (CI), and a lab or digital twin within your operating models. You're likely looking for better integration with your coding or DevOps harnesses today, as builders, to give you intent-based, declarative control over an at-scale but more homogeneous environment than carriers or enterprises face. For many of you, KPIs will shift from uptime to GPU utilization as AI training and inference workloads start to dominate. Start with the sections on deterministic execution and check out the Network Automation Forum's early framework in the Toolkit, read the Trust section, and read the All Domains and Data Centers recommendations.

Feedback is always welcome at research@avidthink.com.

The Wider Current: What Coding Agents and Desktop Copilots Taught Networking

Three Generations of AI Assistants

Predictive AI and other deep learning methods supported network operations and optimization before the current generative and agentic AI era, and they continue to evolve in enterprise, data center, and carrier networks. Meanwhile, generative and agentic AI progress in network AIOps tracks changes and breakthroughs in AI desktop copilots and coding-agent harnesses. It's therefore instructive to take a quick look at where we are with AI desktops and harnesses.

The first generation (circa 2023–2024) was the chat assistant and the software integrated development environment autocomplete; its networking equivalent was the conversational assistant covered in our 2024 report. The second (circa 2025) was the coding agent (Claude Code, OpenAI's Codex, Cursor, Gemini CLI, AWS's Kiro), which could take a goal, plan, edit, run tests, and iterate; our 2025 report cited them as "the art of the possible."

The third (2026) is the desktop or cloud agent that works across a person's files and applications. The control pattern embedded in these harnesses has migrated into networking almost intact: networking vendors now run the same seven control elements, from plan-then-execute to token budgets with model routing. See the figure for how each element maps to its networking counterpart.

©2026 AvidThink LLC. All Rights Reserved.

THE AGENT HARNESS PATTERN, FROM CODING AGENTS TO NETWORK OPERATIONS

2023–2024	2025	2026
Chat assistant, IDE autocomplete	Coding agent: plan, edit, run tests, iterate	Desktop / cloud agent across files and apps
Conversational network assistant (2024 report)	Agentic frameworks – "the art of the possible" (2025 report)	Governed agents acting on networks through a harness

Three generations of AI assistants: the AI industry (teal) and the networking equivalent (purple)

	Coding-agent harness element	Networking equivalent
1	Plan, then check before execution	Proposed change plan and dry run reviewed before any push
2	Permission prompts; allow-lists for routine actions	Approval gates for consequential changes; pre-approved action classes
3	Execution in a sandbox or a branch	Digital twin, lab replica, or GitOps branch before production
4	Context in files (CLAUDE.md, AGENTS.md), not discovered at run time	Pre-built context: source of truth, ontology / knowledge graph, topology
5	Repeatable procedures packaged as "skills"	Runbooks and SOPs packaged as agent skills
6	MCP as the common tool protocol	MCP servers and gateways on controllers, inventory, ticketing
7	Token budgets, bring-your-own models, model routing	Token targets per team, BYO-LLM, small models for the routine path

AvidThink

Source: AvidThink.

avidthink.com

This year, networking vendors are no longer talking about which model (or models) they use; they're showcasing their workflows, frameworks, and harnesses. This mirrors the messaging from enterprise AI software vendors. The common refrain is that today's models (even open-weights or smaller models) are good enough for many tasks. Still, reliability and performance come from the right harness paired with the right context.

Where Networking Diverges from Coding Agents

Coding has something networking and many other enterprise tasks do not: a cheap, automatic verifier. It's easy for a compiler, test suite, and CI pipeline to give a coding agent immediate feedback on what went right or wrong. Verification of changes to networks can be much harder. For example, did the remediation step resolve the intermittent quality-of-service issue a laptop client was seeing on the Wi-Fi network?

Coding has something networking does not: a cheap, automatic verifier. Digital twins and lab replicas partially substitute for the missing verifier.

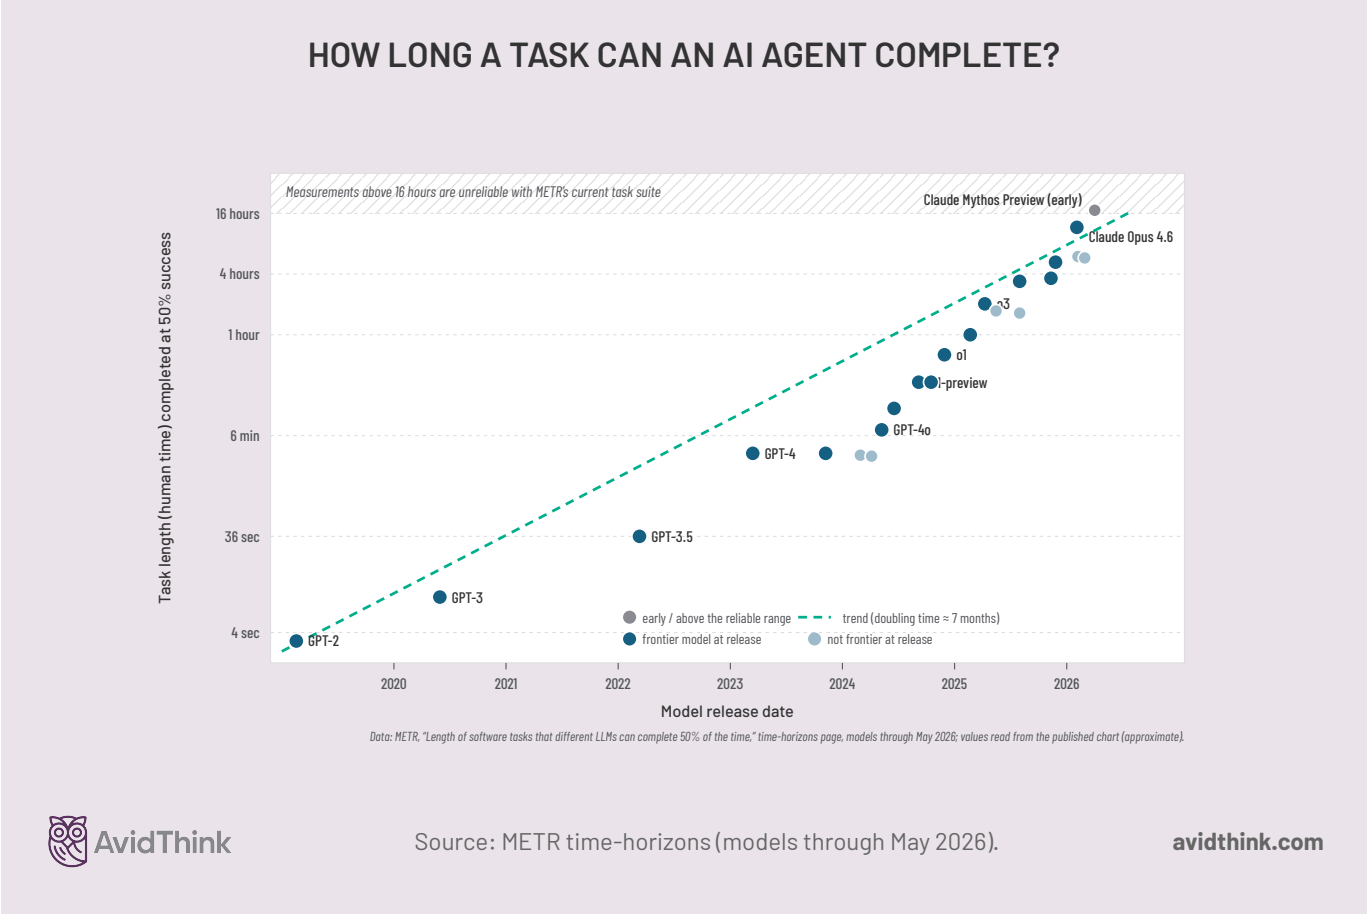
This is where other verification techniques, such as digital twins and lab replicas, can provide strong feedback. While less definitive than feedback in coding, they let an agent's recommended steps be tested for feasibility and impact before production roll-out. Furthermore, the same digital twin can supply topology and recent-change context to AI-assisted root cause analysis.

©2026 AvidThink LLC. All Rights Reserved.

A second difference, related to using digital twins for verification, is that networking AI agents may need context pre-built and curated in advance. Coding and desktop agents try to build context at runtime. However, large networks in enterprises, data centers, and carriers can emit so much telemetry and event data that they overwhelm even frontier models' context windows.[1] Techniques such as pre-filtering, or the use of graph databases to reflect topology, are used to make the problem tractable. Interestingly, as token costs rise, we're also seeing token-saving schemes in coding and other tasks that pre-build context to make enterprise AI agents more token-efficient. So these two domains might soon converge in their approaches.

The Year of Capable and Rogue Agents

With models able to take on tasks for longer durations (see the METR update below)[2] and harnesses such as OpenClaw (launched November 2025) and its brethren, 2026 was the year of personal agents.



At the same time, the industry has yet to solve the alignment problem, and agent governance and control lag behind their capabilities. Since the OpenAI/Hugging Face incident,[3] reports of agents gone "rogue" have multiplied, with higher-value targets. But the word "rogue" is a misnomer. The agents act as they were trained; we lack the techniques to monitor and bound them adequately. For constrained tasks, with guardrails in place and humans in the loop or on the loop, agentic AI is useful and, in single, well-bounded scenarios, may run unattended. The industry (networking and beyond) is active and has huge incentives

[1] V. Fernández (Colt Technology Services), "From Data to Decisions: Real-Time Network Context," AutoCon 5, Munich, Jun. 12, 2026. https://www.youtube.com/watch?v=GvZjayObzWQ
[2] T. Kwa et al., "Measuring AI Ability to Complete Long Tasks," arXiv, 2025. https://arxiv.org/abs/2503.14499
[3] "OpenAI and Hugging Face partner to address security incident during model evaluation," OpenAI, Jul. 21, 2026 (updated Aug. 26, 2026). https://openai.com/index/hugging-face-model-evaluation-security-incident/

©2026 AvidThink LLC. All Rights Reserved.

to solve this. Substantial industry efforts have established agent identity and role-based controls, accountability mechanisms, zero-permission sandboxes such as NVIDIA's OpenShell,[4] and out-of-band observability and governance.

The 2026 AI for Networking Reality Across Domains

Claims about "autonomy" in 2026 are limited in scope. TM Forum has its six-level framework, popular with communications service providers (CSPs) and equally usable by enterprises: five capabilities (awareness, analysis, decision, execution, intent) scored from Level 0 to Level 5, where a Level 4 validation covers "select scenarios" — one operational flow in one domain, such as fault management in the radio access network (RAN).[5] The levels are scored on a maturity rating, but they don't answer the question: what is a system permitted to do?

For that, this report uses a common-sense classification of five operating modes that may align better with enterprise users — we've abstracted these from multiple vendor briefings. We couch these as operational permissions. Networking vendor terms may vary ("trusted actions," "Operator mode," "glass box autonomy"), but they map onto the same concepts; the TM Forum correspondence in parentheses is approximate, because TM Forum's levels score capability rather than permission:

- **Mode 1 — Observe and explain:** Read-only. The system gathers evidence, correlates it, and explains what is happening and why (root cause analysis). No action is taken. (≈ TM Forum L0–L1)
- **Mode 2 — Recommend:** The system proposes specific actions with its reasoning; a person decides and executes. (≈ L1)
- **Mode 3 — Execute with individual approval:** The system prepares and executes the change, but each change waits for a person to approve it. (≈ L1–L2)
- **Mode 4 — Execute approved action classes within limits:** Standing authorization for a bounded set of actions (a class, such as wired-port remediation, or a parameter within a range), with limits set in advance and no per-change approval. (≈ L2–L3)
- **Mode 5 — Operate a defined closed loop under policy and continuous verification:** Unattended. The system senses, decides, and acts within policy, with verification before and after each action and automatic rollback. (≈ L3–L4)

Scope is a separate axis. A mode is granted for one scenario, for a whole domain (core, RAN, transport, a campus, a data center fabric), or across domains. So when a carrier says they've achieved TM Forum Level 4 validation, that would be a Mode 5 permission applied to a single scenario in a single network domain (e.g., RAN, transport). And TM Forum Level 5 would correspond to Mode 5 across domains, with agents arbitrating between them.

Where the industry stands in 2026: enterprise vendor GA products stop at Mode 4; telco operator-built systems have reached Mode 5 for one scenario in a single domain; and no organization that we know of runs at Mode 5 across domains (which also means no one has achieved TM Forum Level 5). Next, let's look at the autonomous network reality inside each domain.

[4] R. Ghose, "How ServiceNow and NVIDIA are changing telecom operations," ServiceNow Workflow, Jun. 23, 2026. https://www.servicenow.com/workflow/industries/changing-telecom-operations-nvidia.html

[5] D. Kaur, "What Level 4 autonomous networks actually means, and what it does not," Telecoms Tech News, Aug. 17, 2026. https://www.telecomstechnews.com/news/level-4-autonomous-networks-explained/

©2026 AvidThink LLC. All Rights Reserved.

Communications Service Providers (Access, Transport, Core, Backbone)

Resetting the autonomous networks timelines. Across major telecom conferences in 2026 — Mobile World Congress (MWC), TM Forum Digital Transformation World (DTW), FutureNet World/Asia — telecom executives shared optimism about the path to autonomous networks, but also reset industry expectations. Select operators that re-ran the TM Forum benchmark on themselves this year came out below their earlier self-assessments — Telstra, which had self-assessed at Level 3 to 3.5, found itself "below 2" across all domains[6] — and the Forum's own CTO now places the industry "between about 1.5 and 2.1, realistically."[7] The declared Level 4s are narrow: NBN Co and Ericsson validated Level 4.0 for one scenario, throughput management on a fixed-wireless service covering 800,000 premises;[8] Telefónica reports twelve operational Level 4 use cases;[9] China Mobile, arguably the leader, claims 21.[10] Operators self-assess and submit evidence to a validation program that checks methodology rather than network performance, and TM Forum's own March 2026 report conceded its evaluation tool is "excessively subjective and high-level."[11]

> **"Across all our domains we are below 2 now… I will encourage every CTO in here to test the benchmark."**
>
> TELSTRA, AFTER RE-RUNNING THE TM FORUM BENCHMARK

Promising progress across multiple use cases. Nevertheless, the industry is making progress. Deutsche Telekom's RAN Guardian, built with Google on Gemini, autonomously triggered more than 100 remediation actions in its first month — agentic, operator-built, one scenario (event-driven RAN).[12] NTT DOCOMO runs agents in production on Amazon Bedrock AgentCore that analyze traffic and alarm data from more than a million devices and recommend actions to maintenance engineers;[13] the agents reason and propose, follow approved standard operating procedures (SOPs), and stage via GitOps, but a human approves.[14] Verizon's more than 70 million closed-loop configuration changes in 2025[15] came from rules-based automation with, in its CTO's words, "no tokens in the loop" — conventional closed loop, not agentic. As we'll see in our architecture update below, deterministic automation as the action layer is a recurring design pattern. Verizon's stated goal is to move beyond rules to Level 4.[16]

[6] K. Krogh Andersen (Telstra), "Building the AI-Native Telco – Re-Architecting the Network from the Ground Up," FutureNet Asia 2026, Singapore, Sep. 22, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[7] J. Blackman, "The road to Level 4 – from point solutions to wall-to-wall autonomy," RCR Wireless News, May 19, 2026. https://www.rcrwireless.com/20260519/carriers/road-to-level-4-point-solutions-wall-to-wall-autonomy-tm-forum

[8] "NBN Co and Ericsson achieve Australia-first TM Forum Level 4 Autonomy Validation," Ericsson, Jun. 25, 2026. https://www.ericsson.com/en/press-releases/7/2026/nbn-co-and-ericsson-achieve-australia-first-tm-forum-level-4-autonomy-validation

[9] "Telefónica accelerates its transformation to Autonomous Networks and reaches 12 Level 4 use cases," Telefónica, Feb. 18, 2026. https://www.telefonica.com/en/communication-room/press-room/telefonica-accelerates-transformation-autonomous-networks-reaches-12-level-4-cases/

[10] "DTW 2026 | China Mobile EVP Li Huidi: Autonomous Networks in Future, Faster," Developing Telecoms (Huawei partner content), Jun. 24, 2026. https://developingtelecoms.com/telecom-technology/wireless-networks/20432-dtw-2026-china-mobile-evp-li-huidi-autonomous-networks-in-future-faster.html

[11] M. Wagner, "Telcos pick up the pace to Level 4 autonomous networks, says TM Forum," Fierce Network, Apr. 3, 2026. https://www.fierce-network.com/cloud/telcos-hit-level-4-autonomous-network-milestone-says-tm-forum

[12] "Deutsche Telekom and Google Cloud Collaborate for Superior Network Experience with Agentic AI," Deutsche Telekom, Feb. 25, 2026. https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/2/mindr-ai-agents-in-telekom-network

[13] "DOCOMO Begins Commercial Deployment of Agentic AI System for Network Maintenance Using One of the World's Largest Datasets," NTT DOCOMO, Feb. 25, 2026. https://www.docomo.ne.jp/english/info/media_center/pr/2026/0225_01.html

[14] Y. Nakajima (NTT DOCOMO), "NTT DOCOMO's Case Study: Agentic AI-Powered Autonomous Network from Day 0 to Day 2 for Cloud-Native Functions," FutureNet Asia 2026, Singapore, Sep. 22, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[15] Y. Tenorio, "Architecting autonomy: Network infrastructure for the agentic era," Verizon, Jun. 24, 2026. https://www.verizon.com/about/news/verizon-architecting-network-autonomy

[16] M. Wagner, "Verizon: 'We are pushing hard toward Level 4 autonomy in critical segments'," Fierce Network, Jun. 24, 2026. https://www.fierce-network.com/cloud/verizon-we-are-pushing-hard-toward-level-4-autonomy-critical-segments

©2026 AvidThink LLC. All Rights Reserved.

Enterprise Campus and WAN

Enterprise networking vendors are shipping solutions with similar but evolving AI architectures. Most incumbent networking vendors now offer conversational assistants, and use is growing (Cisco's handles about 150,000 queries a month).[17]

Generative dashboards (use of AI to create personalized dashboards with custom widgets reflecting important KPIs) run the gamut. And in 2026, **AI-assisted actions** are the next step up. However, every vendor's GA product still depends on a human in the loop (HITL) for anything not pre-authorized. HPE's Marvis self-driving actions execute operator-enabled action categories ("trusted actions" such as wired-port remediation) across both Aruba Central and Mist.[18] Cisco's Agentic Actions is in beta, with changes requiring approval.[17] Extreme's Agent ONE Coworker, GA since September 2026,[19] uses proactive Nudges to flag emerging issues, investigates likely causes, and recommends actions for a person to approve (Mode 3);[20] Operator mode (Modes 4–5) is expected in the first half of 2027.[21] Arista's AVA (Autonomous Virtual Assist) requires explicit administrator approval before corrective changes.[22]

However, it's harder to pin down how many enterprises use these features in production than the vendor usage counts suggest. Even vendors, when pressed, acknowledge that many features are new and adoption stats reflect growth from a tiny baseline. The one positive advancement is that all major vendors have some form of impairment detection and troubleshooting, with suggested remediation actions. Many network engineers tell us this common capability has cut their mean time to repair/recovery (MTTR) by sharply reducing troubleshooting time. This is anecdotal rather than measured, but it shows real-world value.

Data Center Operations

The data center domain exhibits different dynamics. The pattern that's emerging is a dev-centric framework built around GitOps, CI, and a digital twin or lab for verification before committing a change (for large-scale or multi-site deployments). As AI workloads dominate, we expect KPIs to shift from uptime to GPU utilization and the network metrics behind it (job completion time, network fabric congestion). After all, networking-for-AI spending outruns AI-for-networking spending, and will for years.

One interesting example came from a vendor's own IT shop: Nokia IT moved six of its 21 data centers to SR Linux and Event-Driven Automation (EDA) in about a year, with intent captured as YAML in GitLab, an on-demand digital twin, and Cursor as its engineers' front end.[23] Outages fell from two or three a month to one 40-minute incident in ten months, roughly a twenty-fold reduction.[24] In a different example, HPE Services used HPE Apstra's digital twin and cabling checks to bring up and validate an 8,000-link, 4,000-GPU cluster at Oak Ridge National Laboratory in less than five weeks. But this was an automation deployment without agentic AI.[25] Both are evidence that deterministic automation, good network operations (NetOps) practices, and a good source of truth can deliver strong operational gains. Any deployment of reasoning agents and agentic AI should leverage proven underlying automation capabilities.

[17] S. Goel and K. Shah (Cisco), "Cisco Secure Networking: Intelligent Autonomous Action with AgenticOps," Networking Field Day 41, Oct. 9, 2026. https://techfieldday.com/video/cisco-secure-networking-intelligent-autonomous-action-with-agenticops/

[18] "HPE expands self-driving networks across edge, campus, data center, and AI factories," Hewlett Packard Enterprise, Jun. 16, 2026. https://www.hpe.com/us/en/newsroom/press-release/2026/06/hpe-expands-self-driving-networks-across-edge-campus-data-center-and-ai-factories.html

[19] "Extreme Delivers Industry's First Proactive, Context-Aware AI Agent for Networking," Extreme Networks, Sep. 9, 2026. https://investor.extremenetworks.com/news/news-details/2026/Extreme-Delivers-Industrys-First-Proactive-Context-Aware-AI-Agent-for-Networking/default.aspx

[20] M. Wagner, "Extreme Networks moves beyond AI chatbots for network operations," Fierce Network, May 6, 2026. https://www.fierce-network.com/broadband/extreme-networks-moves-beyond-ai-chatbots-network-operations

[21] "Extreme Agent ONE Coworker Saves Customers 1,000 Hours in First Month of Availability," Extreme Networks, Oct. 21, 2026. https://s205.q4cdn.com/756692490/files/doc_news/Extreme-Agent-ONE-Coworker-Saves-Customers-1000-Hours-in-First-Month-of-Availability-2026.pdf

[22] K. Duda and S. Ventkiteswaran (Arista), "Arista AVA – Agentic AIOps Platform & Demo," Mobility Field Day 14, May 8, 2026. https://techfieldday.com/video/arista-ava-agentic-aiops-platform-demo/

[23] P. Pandey and M. Bushong (Nokia), "Modernizing IT Networks with Nokia EDA and SR Linux," Networking Field Day 41, Oct. 7, 2026. https://techfieldday.com/video/modernizing-it-networks-with-nokia-eda-and-sr-linux/

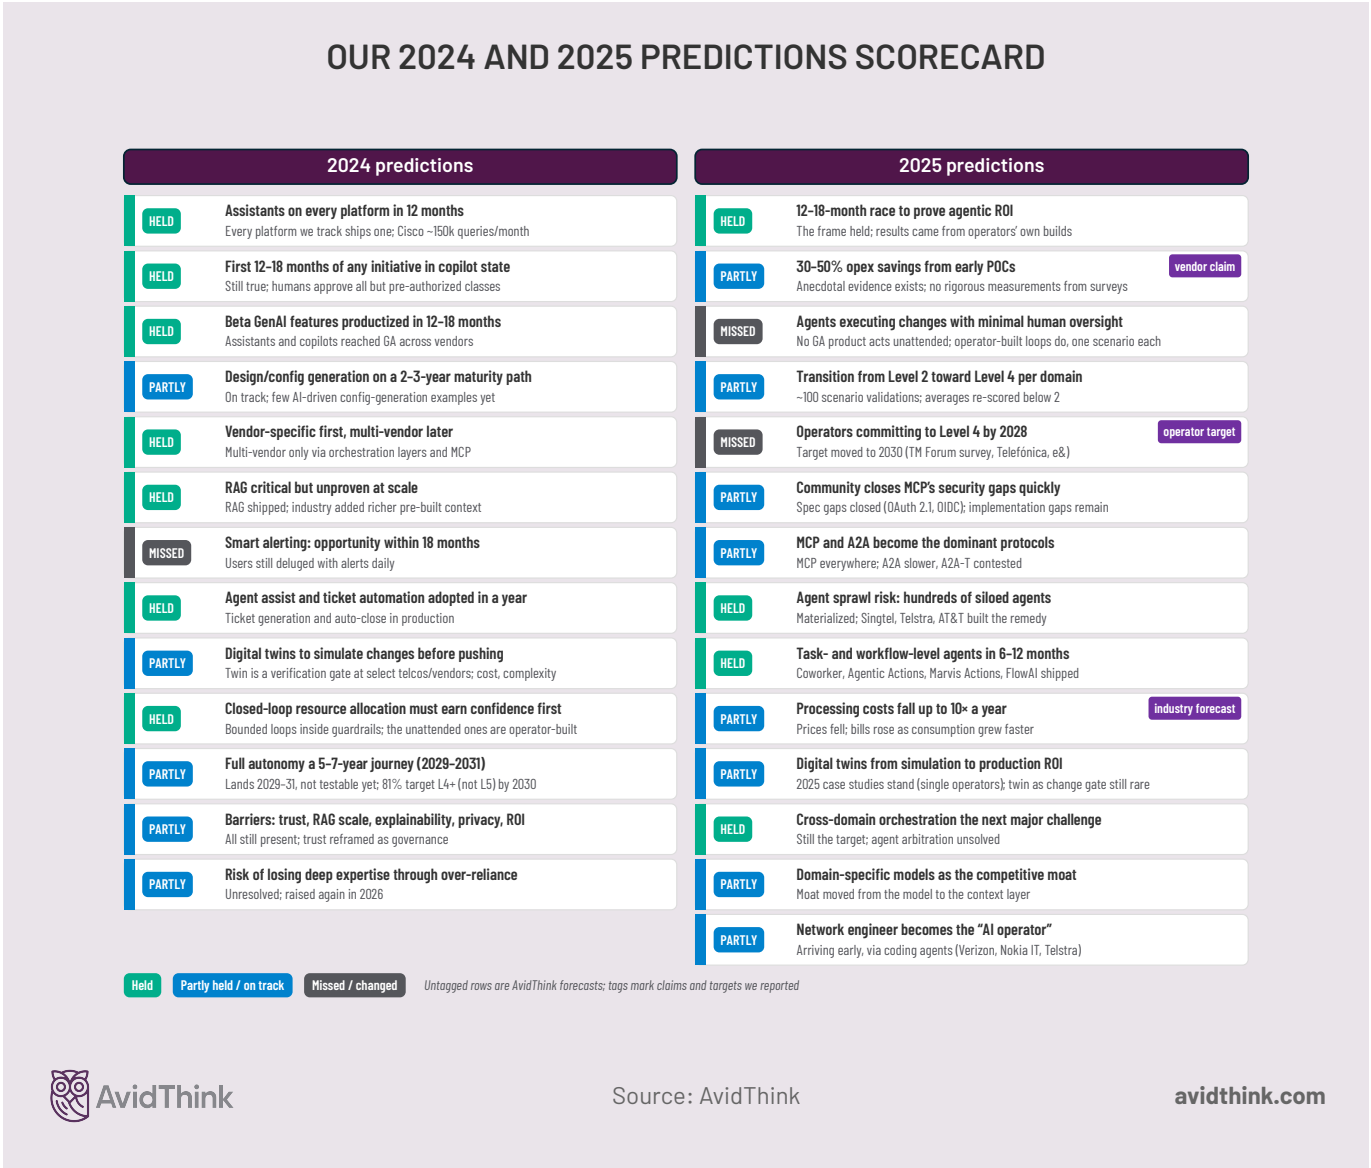
[24] A. Kumar and A. Abutaleb (Nokia), "Nokia Network Coding and Testing – Demo of What We Do Everyday," Networking Field Day 41, Oct. 7, 2026. https://techfieldday.com/video/nokia-network-coding-and-testing-demo-of-what-we-do-everyday/

[25] S. Narahari (HPE), "HPE Networking Data Center: AI Data Center Operations in Action," Networking Field Day 41, Oct. 8, 2026. https://techfieldday.com/video/hpe-networking-data-center-ai-data-center-operations-in-action/

©2026 AvidThink LLC. All Rights Reserved.

What Changed in 2026: Our Scorecard

We made predictions in each of our previous two editions. Here, we'll take a quick pause in the report to grade ourselves. In scoring, we separate our own forecasts from the vendor claims and operator targets we reported at the time. Our scorecard, expressed visually:



Reviewing Our 2024 Predictions

Most of what we wrote in 2024 held, and our timing predictions weren't far off. We predicted assistants on every major platform within 12 months, beta GenAI features productized within 12–18 months, and vendor-specific solutions before multi-vendor ones arrived. Those were on the nose. We also called RAG critical but not scalable on its own. One of our predictions lasted longer than we expected: the 12–18-month copilot state. We're still in that state.

©2026 AvidThink LLC. All Rights Reserved.

Meanwhile, our 5–7-year horizon to full autonomy lands in 2029–2031, so we still have time to get it right (or wrong). Carriers in the industry concur (for now): 81% of 80 operators in TM Forum's survey target Level 4 or above by 2030.[26] We expect that as each telco gets into implementation details and measures real-world outcomes, it will find that moving beyond Level 3 takes more effort than planned.

We fell short on smart alerting (enterprise users are still deluged daily) and digital twin adoption for simulating changes (only select telcos and enterprise vendors offer it because of cost and complexity). Every barrier we listed in 2024 remains, and the risk of losing deep networking expertise through over-reliance on AI is unresolved. If junior network engineers don't move up the ranks and develop the intuition and judgment to become senior engineers, who will eventually supervise the agents or save the network when AI fails?

Tracking Our 2025 Predictions

The 2025 edition predicted a race among the major platform players to prove the ROI of their agentic frameworks, and that became the framing through the year and into 2026. But instead of strong ROI figures published by vendors, most of the results we are seeing come from end users: carriers and vendors' own IT departments (e.g., Telstra, AIS, DOCOMO, Nokia IT, Cisco IT). Today, many agentic frameworks are still running in proof-of-concept mode across enterprises and telcos.

We also predicted that the community would close MCP's security gaps. The MCP specification has closed these gaps (the Interoperability section covers the 2026 revision), but implementation lags. Therefore, we expect continued MCP-related security incidents in 2026–2027. MCP is indeed everywhere, as predicted, but A2A uptake is slow in networking, and its telco profile (A2A-T) isn't widely embraced. The agent-sprawl risk we flagged materialized, and the leading network operators who aggressively rolled out agents are grappling with it.

Our misses include overstating how soon agentic systems that could "execute configuration changes with minimal human oversight" would arrive as products. The 30–50% operational cost savings we reported have yet to be widely achieved (or measured). "Level 4 by 2028" was an operator target, declared from conference stages and reported by us, and it has slipped to 2030 in the TM Forum's latest survey and in carriers' own statements.

How much generative or agentic AI is in production? It varies.

Four 2025–2026 surveys give answers from 3% to 51%, depending on what was counted. The Network Automation Forum (NAF) counts large language models (LLMs) with access to production systems (3% of the 526 who answered that question);[27] GSMA Intelligence finds more than two-thirds of operators have not progressed beyond Level 2;[28] 21% of respondents to TM Forum's survey report Level 3 or above;[11] and Cisco's Omdia-conducted survey counts leaders who say their agentic AI "acts in production today" (51%).[29] All four can be true at once. Cisco is a vendor, so we would expect its number to be the most optimistic; the NAF figure is the strictest and, for enterprises, likely closest to ground truth.

AI processing costs fell as we predicted, but bills rose because token consumption rose faster with widespread use of expensive reasoning models. Finally, digital twins: the 2025 edition cited single-operator, vendor-reported results (Chunghwa Telecom's

[26] "Autonomous Networks Leadership Forum spotlights Level 4 shift from cost-cutting to revenue growth," TM Forum, Jun. 2026. https://www.tmforum.org/news-insight/newsroom/autonomous-networks-leadership-forum-spotlights-level-4-shift

[27] C. Grundemann, "The 2025 State of Network Automation: What's Really Happening in Production Networks," Network Automation Forum, presented at NANOG 95, Oct. 28, 2025 (slides). https://storage.googleapis.com/site-media-prod/meetings/NANOG95/5509/20251025_Grundemann_The_2025_State_v1.pdf

[28] B. Kizhakedath, "GSMA Intelligence: Telecom Operators Must Accelerate Shift to Intent-Driven Autonomous Networks," TelecomLead, Mar. 11, 2026. https://telecomlead.com/5g/gsma-intelligence-telecom-operators-must-accelerate-shift-to-intent-driven-autonomous-networks-125105

[29] "Cisco AI Research: AgenticOps Scaling Quickly in the Enterprise," Cisco Newsroom, Sep. 23, 2026. https://newsroom.cisco.com/content/r/newsroom/en/us/a/y2026/m09/cisco-ai-research-agenticops-scaling-quickly-in-the-enterprise.html

©2026 AvidThink LLC. All Rights Reserved.

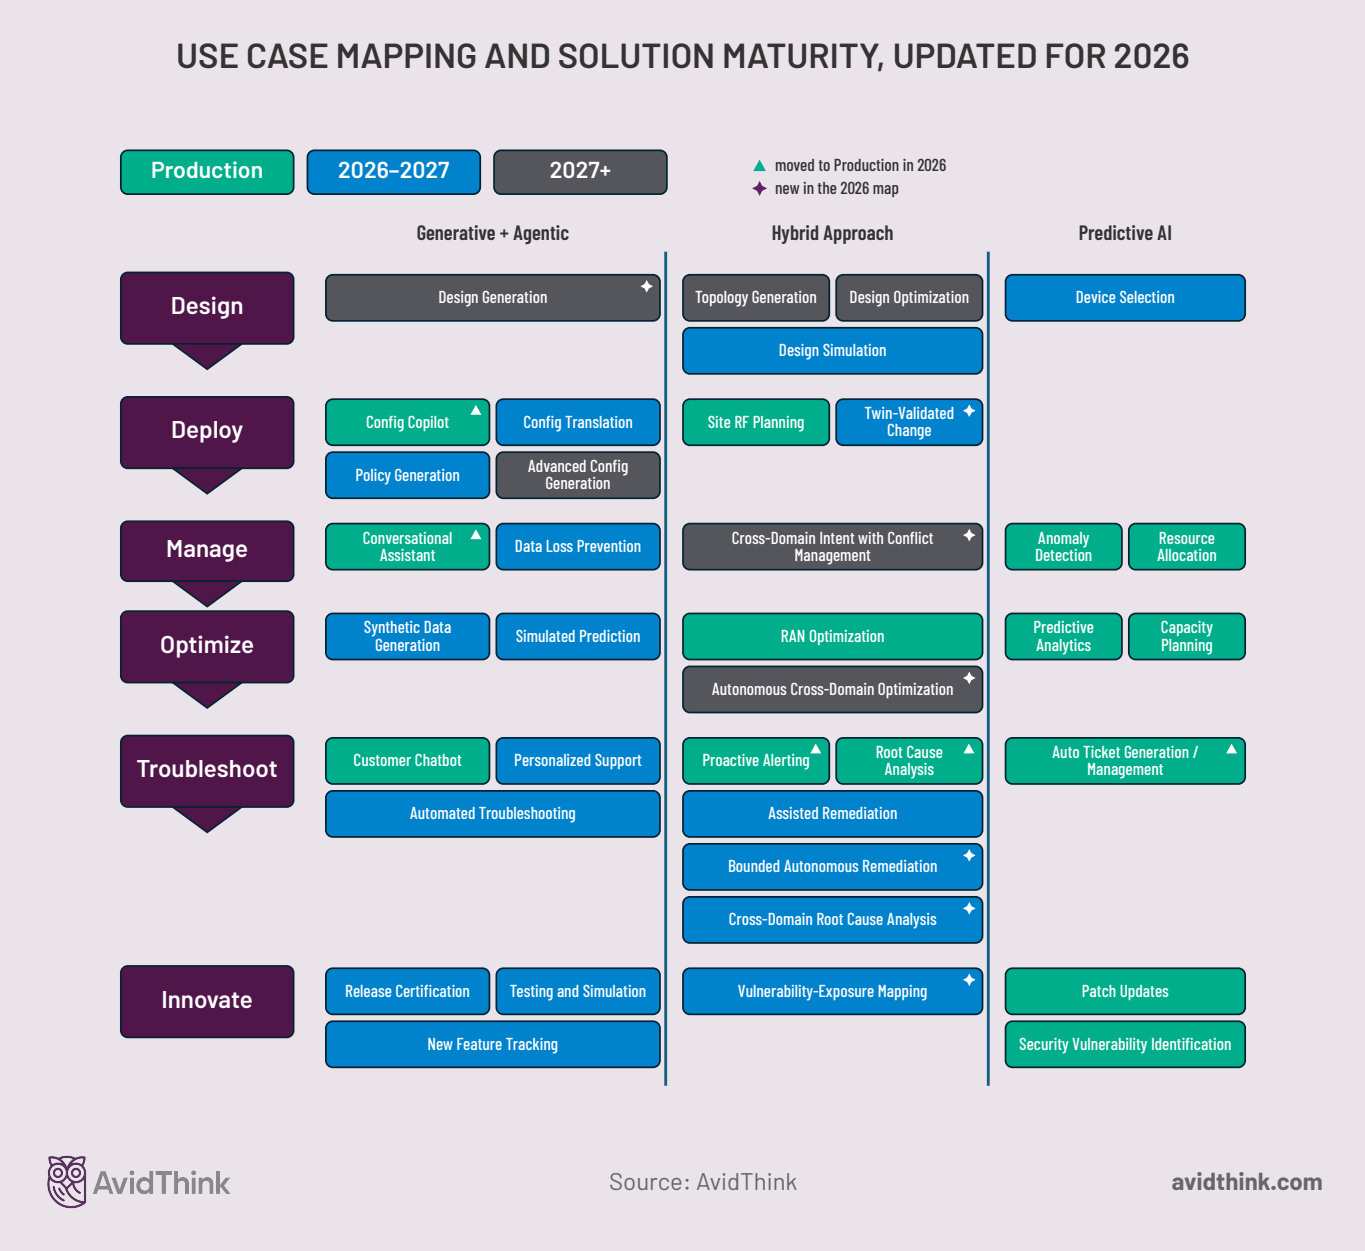
capacity gain, Nokia's shorter deployment cycles), and those results stand. What has not happened is broad adoption of the twin as the verification gate for agent-initiated change.

The 2026 Toolkit: Predictive, Generative, Agentic, and the Harness

Our 2024 taxonomy of AI capabilities and our 2025 three-layer toolkit are still valid. We encourage readers who have not yet read our 2024 or 2025 editions to download a copy. Here's this year's update to our taxonomy and frameworks.

Updated Use Case Map: Vulnerability Exposure Is a Key New Addition

Several items moved into production this year: the conversational assistant, the configuration copilot, root-cause analysis, automatic ticket generation, and proactive alerting.



©2026 AvidThink LLC. All Rights Reserved.

A new entry, vulnerability-exposure mapping from Common Vulnerabilities and Exposures (CVE) feeds to affected devices to compensating controls, is a key candidate (especially given the wider impact of AI-assisted vulnerability exploitation) and enters the map in the 2026–2027 column. Gluware unveiled its Titan Exposure Management product in May 2026 and reports strong enterprise interest, but it is one vendor's new product so far. The 2026–2027 column also includes bounded autonomous remediation (pre-authorized action classes in products, single-scenario loops beyond the operator builds that have them today), configuration translation, policy generation, twin-validated change, and cross-domain root-cause analysis. Cross-domain intent with conflict management, design generation, and autonomous cross-domain optimization will remain in 2027 and beyond, which fits the 2–3-year schedule we predicted in 2024.

Predictive AI: Still the Foundation

Predictive AI remains important for optimization and capacity management, and carriers (radio access, transport, and core networks) and enterprises (Wi-Fi and wired networks) will continue to use it. One emerging pattern is predictive AI working in tandem with generative/agentic AI: a time-series model pre-processes data so the reasoning model isn't overwhelmed by telemetry (and token costs stay under control). IBM's Tiny Time Mixer and TSPulse models, with a few million parameters and GPU-free inference, are built for exactly that front-end role.[30]

Generative AI: From Retrieval Augmentation to Context Engineering

Our 2025 story was about going "beyond RAG." In 2026, RAG still pulls in knowledge bases and operational procedures, but the context is more sophisticated. In some cases, vendors fine-tune existing open-weights models with domain-specific data to embed more knowledge into LLMs or small language models. Colt, NTT, and other telcos also use continuously updated graph databases to provide context.[1] Domain-specific models such as Cisco's Deep Network Model, HPE's Marvis Large Experience Model (which is not an LLM), or SoftBank's telecom models provide improved performance in the domains they target.

Another element that emerged in 2026 is the growing importance of model sovereignty and control—especially outside the United States and for large enterprises. Certain customers want the flexibility to use their own locally hosted models (or models hosted on their private infrastructure). Gluware's Titan is bring-your-own-LLM;[31] Nokia's fixed-access agents let operators choose the LLM;[32] and Ericsson launched a self-hosted NetCloud Private, with global availability in October 2026.[33]

Agentic AI: Harnessing the Value in 2026

The biggest 2026 update in agentic AI is in solution architecture. Vendors are using mostly similar frameworks, regardless of the domains they serve (enterprise, carrier, data center).

A distinct context layer (source of truth, ontology, digital twin) below the agents is key to their accuracy and performance. The context layer taps into the network, filtering and shaping relevant data. Data includes telemetry, configuration, logs, operational manuals, knowledge bases, support ticket databases, and more.

In parallel, a governance plane (identity, registry, gateway, policy, audit) sits alongside the agents, gating and watching their actions. This is necessary to achieve trusted autonomy over time. And the MLOps/Model management component now includes managing token budgets (through model selection and routing) without sacrificing performance and quality.

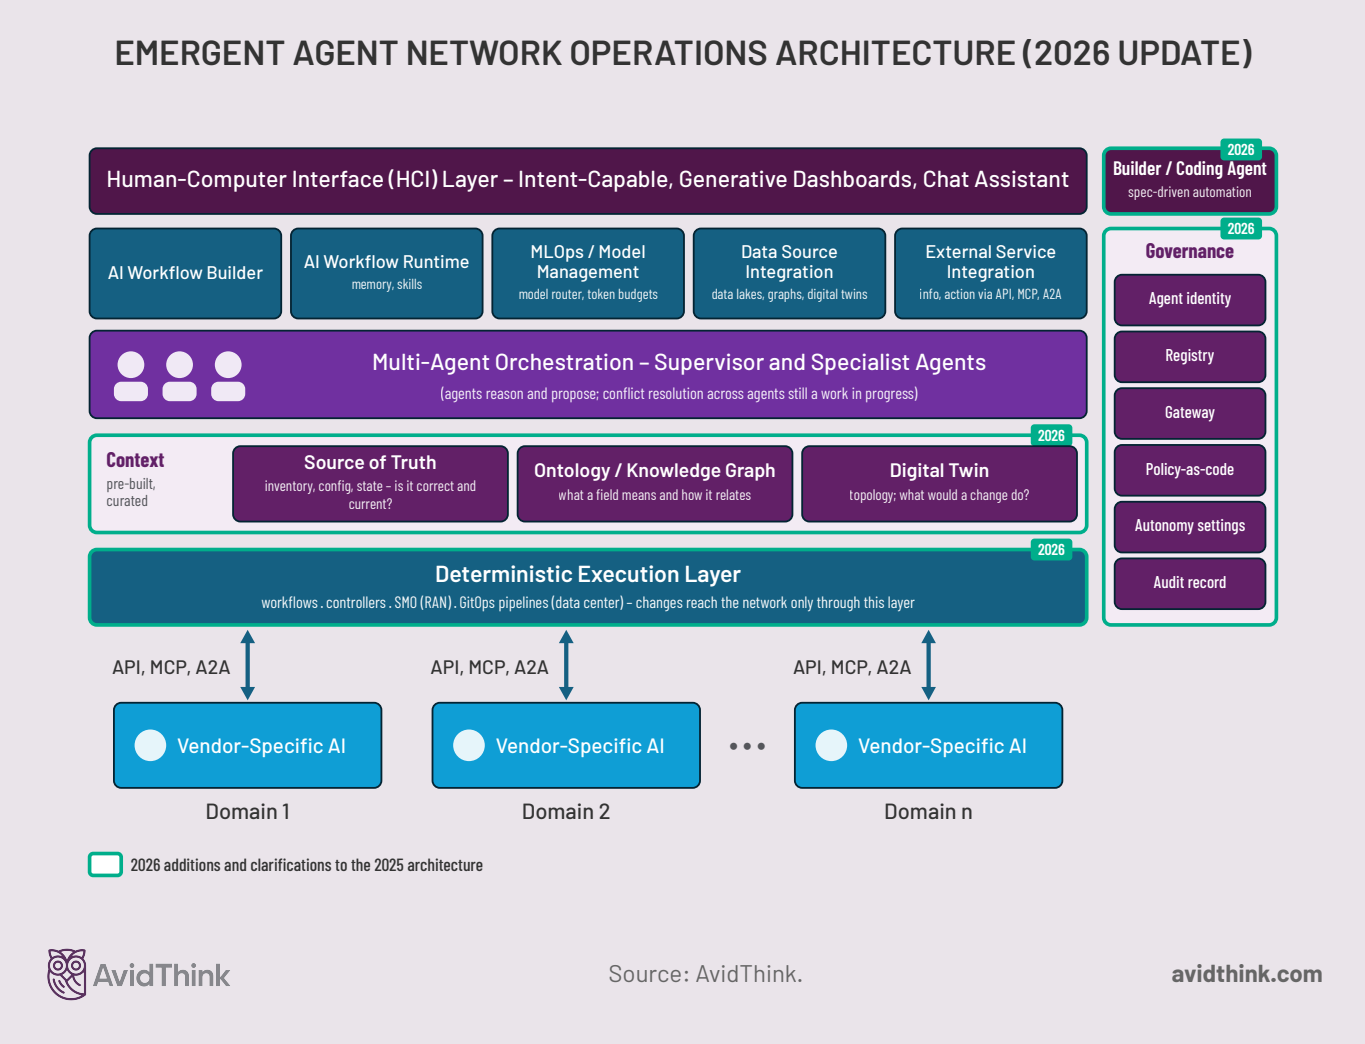
[30] "Granite Time Series," IBM Granite documentation, IBM. https://www.ibm.com/granite/docs/models/time-series/
[31] M. Haugh, "Gluware 5.8: Advancing Intelligent Network Automation with AI, Scale, and Multi-Vendor Innovation," Gluware, Feb. 11, 2026. https://gluware.com/gluware-5-8-advancing-intelligent-network-automation-with-ai-scale-and-multi-vendor-innovation/
[32] "Nokia launches agentic AI for home and broadband networks," Nokia, May 12, 2026. https://www.nokia.com/about-us/news/releases/2026/05/12/nokia-launches-agentic-ai-for-home-and-broadband-networks/
[33] "Ericsson meets Gulf demand with NetCloud Private for customer control in Wireless WAN operations," Ericsson, Sep. 23, 2026. https://www.ericsson.com/en/press-releases/5/2026/ericsson-netcloud-private-boosts-gulf-wireless-wan

©2026 AvidThink LLC. All Rights Reserved.

For those familiar with the evolution of coding agents and agentic clients, this architecture will look familiar. Additions that parallel those enterprise agents include skills (packaged procedures), memory (short- and long-term), and tools accessed over MCP or called via APIs.

As multi-agent coding (and agentic desktops) became more widespread this year, networking vendors adopted the same pattern. Orchestrated by a more capable agent, sub-agents gather the necessary context from a wide variety of historical and real-time data (telemetry and state from network elements, and the other data elements described earlier).

Commercial examples include Extreme's AIOps evolution from a chatbot-centric first generation to an agentic architecture that combines unified network data, a knowledge graph, multi-model orchestration, and specialized skills, similar to agentic workspaces such as Claude Cowork and OpenAI's Codex.



Similarly, Ribbon and AWS describe the pattern in their joint solution: specialist RAN, transport, packet-core, and voice-core agents run in parallel, and a reasoning agent synthesizes their findings into one root cause.[34] Nokia's Gemini-built Assurance Center agents (a router plus five specialists, two fully functional today)[35] and Blue Planet's supervisor-and-specialists demon-

[34] R. Ramanathan (Ribbon) and C. McKenna (AWS), interview in this report.

[35] "Nokia and Google Cloud Partner to Embed AI Agents, Built with Google's Gemini Models, Into Nokia's Autonomous Network Product Suite," Google Cloud, Jun. 22, 2026. https://www.googlecloudpresscorner.com/2026-06-22-Nokia-and-Google-Cloud-Partner-to-Embed-AI-Agents,-Built-with-Googles-Gemini-Models,-Into-Nokias-Autonomous-Network-Product-Suite

©2026 AvidThink LLC. All Rights Reserved.

stration with AWS, with the equipment-modeling agent running on Amazon Bedrock,[36] follow the same architecture.

One observation from 2026 is the architecture's separation between reasoning and acting. In many solutions, the "acting" portion consists of a **deterministic execution layer**. LLM-powered agents reason and propose, but changes reach the network only through workflows, controllers, or pipelines whose behavior is predictable and testable. For example, vendor Itential's FlowAI agents act only through governed workflows; Gluware's sit on DIAL, its device interaction layer; IBM Concert Workflows are triggered as deterministic APIs to separate probabilistic judgment from fixed execution; Amdocs separates agent reasoning and governance from execution through established deterministic workflows. This action layer largely reflects the automation the industry has built over the past decade: products such as IBM/Red Hat's Ansible, Network to Code's Nautobot apps, and orchestration vendors' workflow engines. This layer now serves the reasoning layer above, which can translate "intent" into actual network actions and bring the vision of intent-based networking into production.

Token Economics Becomes Budget Economics

The MLOps/model management layer in the architecture above is where the 2026 cost story plays out. Our 2025 recommendation was "don't overlook the cost," and we predicted processing costs would fall by up to 10× a year. Per-token prices did fall as technologies improved and open-weights models proliferated. But bills rose anyway, because consumption rose faster.

Telstra spends about 15 billion tokens a day, up 275% in a year, across more than 400 use cases. Per-token prices fell; bills rose anyway.

As an example, Telstra spends about 15 billion tokens a day, up 275% in a year, across more than 400 use cases, 12% of them in autonomous networks.[6] Verizon gives each role a token allowance and tiers models by task.[37] On the vendor side, an Amdocs-commissioned Appledore note records token consumption "substantially higher than early projections."[38] Across enterprises and carriers, executives are forced to work within token budgets, and we expect the day will come when they are optimizing for a combined token-plus-staff-salary budget.

At the same time, networking vendors are split on who pays for tokens directly. Cisco's AI is "all you can eat," with anti-abuse caps and limited visibility into token use.[17] Extreme includes Coworker in the Platform ONE subscription.[19] Gluware is bring-your-own-LLM with no metering beyond the per-device fee.[39] Regardless, briefings with vendors indicate all of them are sensitive to conserving tokens, and increasingly they are relying on smaller models for routine operations, and using frontier models for the exceptions or advanced reasoning.[40]

Open Questions in Agentic Network Operations

As these agentic solutions gain traction, unresolved issues remain. Multi-agent orchestration and conflict resolution between agents (when they have competing actions) remain work in progress, especially when different vendors control different domains. For example, if a cost-control agent demands that radios power down but a quality-of-service agent wants improved coverage and power, a higher-level mechanism needs to resolve this. Open architectures and standards forums will likely play a role here in allowing cross-vendor interactions.

[36] S. Verdi, "Blue Planet extends service automation with as-a-service agentic AI scaling with AWS at MWC 2026," Blue Planet, Feb. 26, 2026. https://www.blueplanet.com/blog/2026/blue-planet-extends-service-automation-with-as-a-service-agentic-ai-scaling-together-with-aws-at-mwc-2026
[37] L. Dignan, "Verizon aims for autonomous network: Takeaways on architecture, AI agent returns, managing costs," Constellation Research, Jul. 2, 2026. https://www.constellationr.com/insights/news/verizon-aims-autonomous-network-takeaways-architecture-ai-agent-returns-managing
[38] P. Kelly and R. Curran, "Amdocs Autonomous Network Operations – aOS and the Agentic Telco," Appledore Research (Amdocs-sponsored), Apr. 23, 2026. https://www.amdocs.com/sites/default/files/2026-05/Amdocs-autonomous-network-operations-appledore-research-note-04-26.pdf
[39] "Gluware 5.8: Advancing Intelligent Network Automation with AI, Scale, and Multi-Vendor Innovation," Gluware, Feb. 11, 2026. https://gluware.com/gluware-5-8-advancing-intelligent-network-automation-with-ai-scale-and-multi-vendor-innovation/
[40] C. Wade (Itential), "How NetOps Can Embrace AI to Lead Enterprise Innovation," closing keynote, AutoCon 5, Munich, Jun. 12, 2026. https://www.youtube.com/watch?v=KmXwu7354os

©2026 AvidThink LLC. All Rights Reserved.

Related to conflict resolution is the ongoing NetOps-versus-SecOps (security operations) turf war (primarily within enterprises but also relevant to other domains). If a NetOps agent wants to reroute traffic or modify an ACL to fix a performance issue, how does it negotiate with SecOps guardrails?

Another key question is who owns the context plane: it holds enterprise or carrier proprietary data, such as topology, vendor device inventory, customer-specific policies, and sensitive telemetry. Some context can likely be made portable between customers to improve system performance, but how to use those learnings to fine-tune vendor systems without leaking sensitive customer context and data needs resolution. We'll touch on context more later.

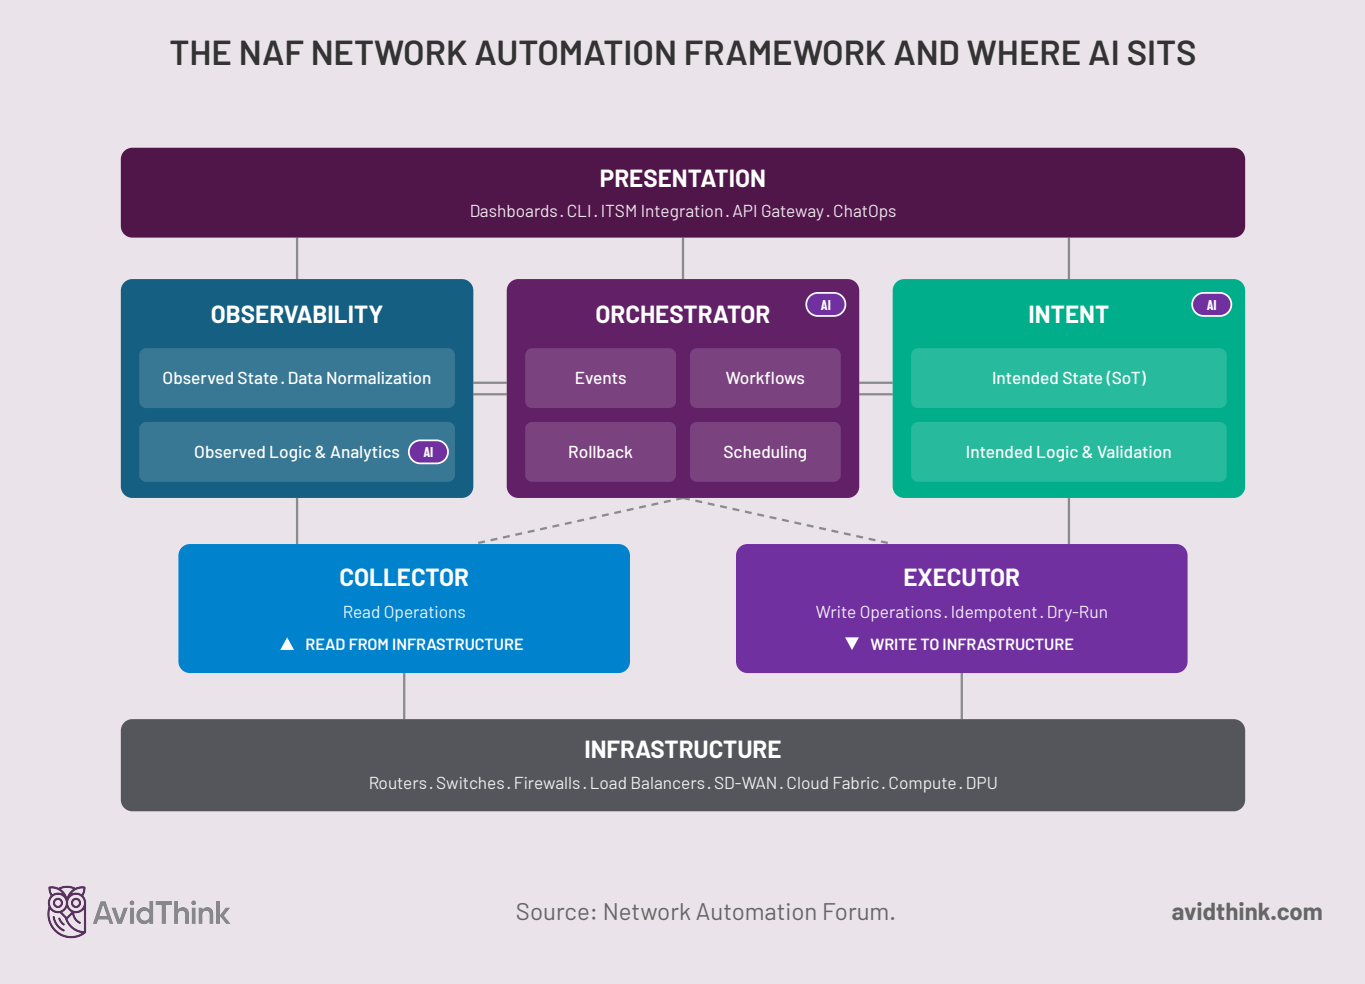
The NAF Framework: A Shared Vocabulary Beneath the Agents

The deterministic execution layer we describe above resembles an emerging reference architecture. The Network Automation Forum, the practitioner community behind the AutoCon conferences, has published its Network Automation Framework as a vendor-neutral reference model built from six functional building blocks: (1) intent (the desired state of the network), (2) observability (the actual state), (3) orchestrator (coordinates workflows across the other blocks), (4) executor (applies changes), (5) collector (reads state over SNMP, syslog, and flow telemetry), and (6) presentation (dashboards, ITSM tools, and CLIs, with authentication and authorization).[41] Where the TM Forum's AI-native Open Digital Architecture is the telco industry's reference frame for agents, this is the enterprise and practitioner one.[42]

It matters for this report because it provides a community view of where agents sit. The framework places AI in three roles: as (1) a consumer of intent (agents need access to the network design in a versioned, schema-validated store), (2) the analytics sub-block of observability, and (3) one of four orchestration styles, alongside scripts, playbooks, and message-bus-driven event coordination. The document is modest by design: a few pages of RFC-style must/should/may lists per block, no mandated protocols or data models, and the authors use "should" rather than "must" in most places. The framework is an emerging work in progress. Whether it heads toward a Cloud Native Computing Foundation-style certification or stays informal is being debated.[43] Regardless, it bears keeping tabs on into 2027.

[41] Network Automation Forum, "NAF Framework," reference documentation. https://reference.networkautomation.forum/Framework/Framework/
[42] "TM Forum members launch AI-native ODA roadmap for autonomous telecoms," TM Forum, Jun. 2026. https://www.tmforum.org/news-insight/newsroom/tm-forum-members-launch-ai-native-oda-roadmap-for-autonomous-telecoms
[43] C. de Luna (EIA / NAF), "NAF Network Automation Framework Update," AutoCon 5, Munich, Jun. 10, 2026. https://www.youtube.com/watch?v=I_iLjkZ3F9I

©2026 AvidThink LLC. All Rights Reserved.



Context: Data, Source of Truth, Ontology, and Digital Twins

We briefly touched on the importance of context earlier. Here, we take a closer look at the context layer, which has grown in significance in 2026 and is a key topic across all networking domains looking to adopt AI for network operations.

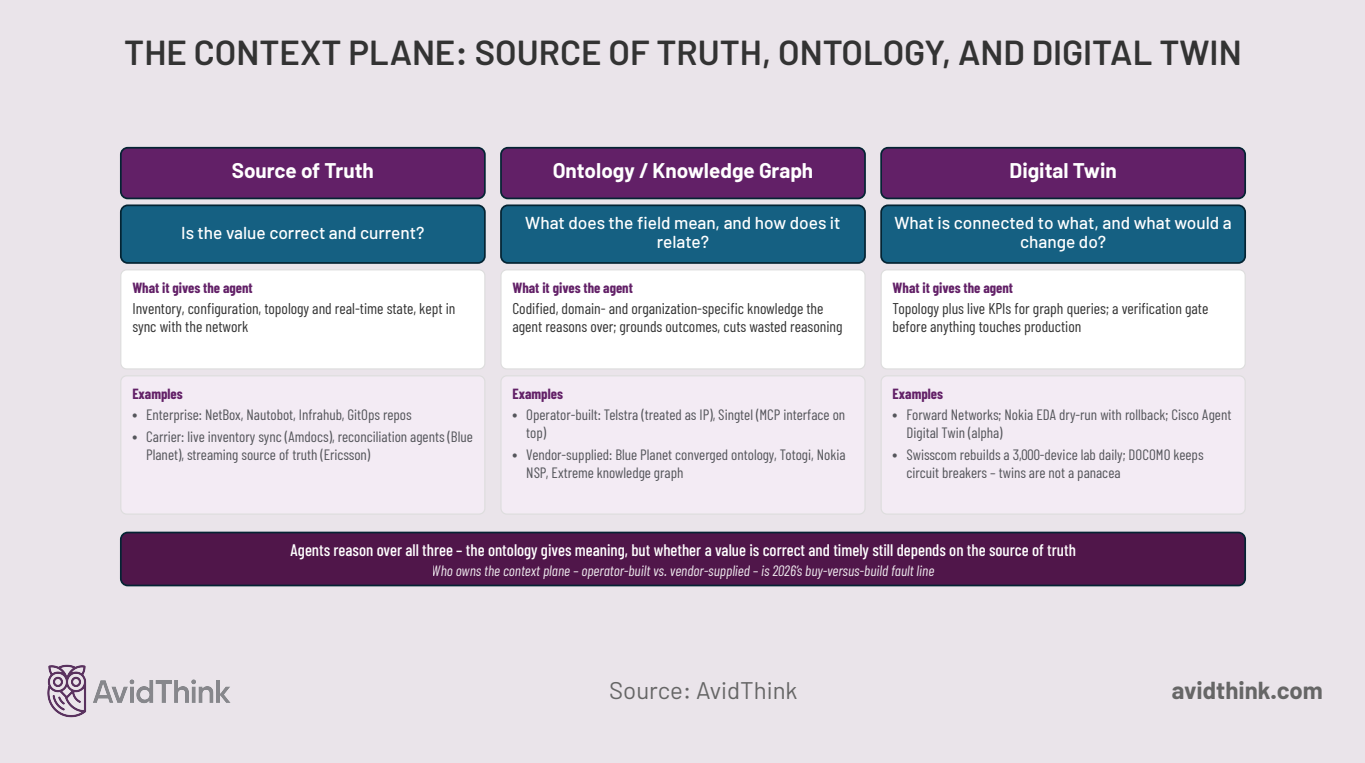
The Data Readiness Problem

In 2026, data availability and readiness remain issues. Data identification, curation, segregation, redaction and anonymization (for privacy and compliance), categorization, and other processing steps remain a burden for enterprises and carriers alike. Network operators have angst about what to save, where to store it, and how long to retain it. Managing myriad databases and data lakes is a huge issue. Especially for telcos, which hold enormous amounts of data across multiple silos, overhead and costs build up, and questions are growing about the usefulness and ROI of retaining copious amounts of data.

For organizations transitioning to more intelligent use of network data (telemetry, events, and so on), one contrarian view we heard this year came from the CTO of Singapore-based mobile virtual network operator MyRepublic, who discarded all legacy data and took a clean-slate approach when he built his AI network operations framework.[44]

[44] L. Chan (MyRepublic), "New Operating Models for the AI-Native Enterprise: Unlocking Value Through Data and Intelligence," FutureNet Asia 2026, Singapore, Sep. 23, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

©2026 AvidThink LLC. All Rights Reserved.



Source of Truth Challenges

Designating a source (or sources) of truth within a network domain that captures both static elements, such as configuration and topology, and real-time events and telemetry remains a challenge.

On the carrier side, advances keep inventory and event databases updated with real-time data. Amdocs uses real-time inventory synchronization and an Operational Digital Twin to reconcile network information for its operational agents;[45] Blue Planet offers a device-reconciliation agent among its specialist agents;[36] and Ericsson added a streaming source of truth for network event data to its automation platform.[46]

In the enterprise, NetBox (a default choice for many organizations), OpsMill's Infrahub, Network to Code's Nautobot, and general GitOps practices are most common. As a reflection of the importance of this category, funds are coming into the space and growth continues: OpsMill raised a $14 million Series A for Infrahub;[47] Selector AI raised $32 million at a $375 million valuation;[48] and Network to Code reported 52% recurring-revenue growth in 2025.[49]

[45] "MWC 2026: Amdocs Unveils CES26, an Agent-driven BSS-OSS-Network Suite, powered by the Amdocs aOS Cognitive Core," Amdocs, Mar. 2, 2026. https://www.accessnewswire.com/newsroom/en/computers-technology-and-internet/mwc-2026-amdocs-unveils-ces26-an-agent-driven-bss-oss-network-sui-1142366

[46] "Ericsson Intelligent Automation Platform expands to support core network automation," Ericsson, Jun. 23, 2026. https://www.ericsson.com/en/press-releases/2026/6/ericsson-intelligent-automation-platform-expands-to-support-core-network-automation

[47] D. Garros, "OpsMill Raises $14 Million in Series A to Fund Product Development & Growth," OpsMill, May 7, 2026. https://opsmill.com/blog/opsmill-raises-14-million-series-a/

[48] "Selector Raises $32 Million to Eliminate Data Silos in Complex Network Operations," Selector, Feb. 18, 2026. https://www.ciodive.com/press-release/20260218-selector-raises-32-million-to-eliminate-data-silos-in-complex-network-oper-1/

[49] "Network to Code Grows Recurring Revenue 52% in 2025, Raises Investment to Accelerate Commercial Automation Software in 2026," Network to Code, Jan. 27, 2026. https://www.accessnewswire.com/newsroom/en/computers-technology-and-internet/network-to-code-grows-recurring-revenue-52-in-2025-raises-investm-1130748

©2026 AvidThink LLC. All Rights Reserved.

Ontology: Critical Information, but Who Owns It?

Every vendor and end-user organization seems to have its own opinion about what an ontology is. Some call it a knowledge layer; some term it a semantic layer. But all agree it is a critical layer that codifies knowledge specific to a network domain and may reflect an organization's unique practices. Agents depend on this layer to provide the critical, relevant information they use for their "reasoning." It grounds agent outcomes and improves system accuracy.

Because of their scale, telecom operators can justify building their own version of this layer and believe they should own it. Telstra views it as intellectual property.[6] Singtel built its own across multiple domains and exposed an MCP interface on top for agents to reduce hallucinations and wasted reasoning steps.[50] Our conversations with executives across many tier-1 carriers worldwide indicate that many of them share the same views on IP ownership of ontology.

Meanwhile, vendors pitch developing and providing that strategic layer as a product differentiator. Blue Planet's ontology stitches data together from planning through orchestration and assurance,[51] so that agents can compare a new situation with how past ones were diagnosed and resolved.[52] Telco software vendor Totogi views its ontology as its differentiation layer. Nokia structures curated Network Services Platform (NSP) data through an ontology.[53] Extreme's networking knowledge graph is a stated competitive moat, grounding AI reasoning in network-specific context.

Who owns the ontology is the buy-versus-build debate of 2026: Telstra and Singtel built their own and call it intellectual property; Blue Planet, Totogi, Nokia, and Extreme pitch theirs as a vendor-provided strategic differentiator.

Nevertheless, even with an ontology, accurate sources of truth remain critical: an ontology gives an agent the semantic meaning of a field, but whether the value is correct and timely depends on a reliable source of truth.

Digital Twins: Context and Verification

Digital twins provide topological context, showing what is connected to what. Digital twins that capture and track the real-time KPIs of their real-world counterparts are valuable targets for graph-based queries from the reasoning layer, checking key metrics as part of a graph walk.

Today, digital twins also serve as verification in many solutions. On the enterprise side, Forward Networks has long sold a twin that simulates what a change would affect across a complex multi-vendor network. Enterprises depend on it for troubleshooting and pre-change simulation. Among the platform vendors, Nokia's EDA pairs an integrated digital twin with dry-run and instant rollback;[54] Cisco's Agent Digital Twin entered alpha in July 2026, allowing AI agents to safely test and validate configuration changes before deploying them to production.[55] Operators build their own, too: Swisscom snapshots its 3,000-router production backbone daily into a scale lab and redeploys everything (about three hours) to catch surprises before production releases.[56] However, digital twins are not a panacea. For instance, Japanese telco NTT DOCOMO keeps circuit breakers in its network

[50] V. Joseph (Singtel), "Data Readiness for AI Deployments at Scale," FutureNet Asia 2026, Singapore, Sep. 23, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[51] K. Anderson (Blue Planet), "Extinction or evolution – why autonomous networks need a new breed of OSS," RCR Wireless News, May 28, 2026. https://www.rcrwireless.com/20260528/opinion/autonomous-networks-need-new-oss-blue-planet

[52] J. Blackman, "Context before control – Blue Planet maps AI architecture for autonomous networks," RCR Wireless News, Aug. 7, 2026. https://rcrwireless.com/20260807/ai/context-control-blue-planet-ai-autonomous-networks

[53] M. Thompson (Nokia), interview in this report.

[54] "Nokia strengthens AI data center performance and AI-enabled automation with enhanced portfolio," Nokia, Nov. 13, 2025. https://www.nokia.com/newsroom/nokia-strengthens-ai-data-center-performance-and-ai-enabled-automation-with-enhanced-portfolio/

[55] M. Cervasio, "Cisco Live 2026: The Agentic Era is Here," WWT, Jun. 11, 2026. https://www.wwt.com/blog/cisco-live-2026-the-agentic-era-is-here

[56] E. Serban and B. Bacal (Swisscom), "Reaching Fully Automated Backbone Networks," AutoCon 5, Munich, Jun. 11, 2026. https://www.youtube.com/watch?v=d4yRszRS6FI

©2026 AvidThink LLC. All Rights Reserved.

because it cannot yet fully rely on the digital twin.[57]

Governance: Policy, Guardrails, Identity, and Audit

Governance Remains Top of Mind

In 2025, we cited trust and control as the foremost barrier. In 2026, that remains true, though the industry is making progress. TM Forum's survey with the IBM Institute for Business Value found that although 72% of CSPs say their AI is trustworthy, only 14% can prove it.[58] The TM Forum DTW Catalyst that won Best Moonshot this year was a governance demo: Zero-Trust Agents for Autonomous Networks (from Telefónica, Orange, Vodafone, Telstra). The team placed a gateway on A2A and MCP calls at every trust boundary and showed live containment of an overreaching and a hijacked agent.[59] And in October, TM Forum and Accenture announced a Trustworthy AI Governance Certification.[60]

Although 72% of CSPs say their AI is trustworthy, only 14% can prove it.

TM FORUM AND IBM INSTITUTE FOR BUSINESS VALUE

Agent Sprawl in 2026: The Forcing Function

In 2025, we warned of "hundreds of siloed, custom-built agents without unified governance." It happened, and it is the forcing function behind the governance plane we describe next. AT&T reported more than 410 generative AI agents in production as early as April 2025.[61] Singtel's "1,000 flowers" approach failed to scale; the group centralized AI under a Chief AI Officer with about 50 use cases.[62] Telstra built an AI control plane and monitors AI like network alerts, in a global operations center.[63]

The Governance Plane in Practice

The components of a governance plane for agents are now more identifiable. The key ones include:

- **Identity.** Each agent needs a verifiable identity of its own, separate from the human who launched it, so that every call between agents and tools can be authenticated and attributed. The TM Forum Moonshot Catalyst mentioned earlier used SPIFFE workload identities with mutual TLS between agents;[59] Cisco's Astrix integration surfaces non-human identities in Duo.[64] The principle is a trust framework between humans and agents and among agents, with clear accountability in the audit trail.
- **Registry and gateway.** A registry records which agents and tools exist and their associated permissions. The gateway sits in the call path to enforce policy and log every MCP, A2A, or API call. For example, Vodafone runs an MCP catalog and

[57] Y. Nakajima (NTT DOCOMO), "Agentic AI & Digital Twins for NetOps: Driving Autonomous Networks," FutureNet Asia 2026, Singapore, Sep. 23, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[58] "TM Forum reveals AI trust gap: 72% of CSPs say their AI is trustworthy, only 14% can prove it," TM Forum, Jun. 2026. https://tmforum.org/news-insight/newsroom/tm-forum-reveals-ai-trust-gap-72-of-csps-say-their-ai-is-trustworthy,-only-14-can-prove-it

[59] T. O'Rourke, "Zero-Trust Agents for Autonomous Networks: winning Best Moonshot Catalyst at DTW Ignite 2026," Solo.io, Jul. 2, 2026. https://www.solo.io/blog/zero-trust-agents-for-autonomous-networks-winning-best-moonshot-catalyst-at-dtw-ignite-2026

[60] A. Wooden, "TM Forum and Accenture launch AI trust initiative," Telecoms.com, Oct. 6, 2026. https://www.telecoms.com/ai/tm-forum-and-accenture-launch-ai-trust-initiative

[61] M. Robuck, "Feature: Operators call up agentic AI," Mobile World Live, Apr. 16, 2025. https://www.mobileworldlive.com/att/feature-operators-call-up-agentic-ai/

[62] K. Leong (Singtel), "Capturing the AI Opportunity: What Will It Take for Telcos to Win in the AI Economy?," FutureNet Asia 2026, Singapore, Sep. 23, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

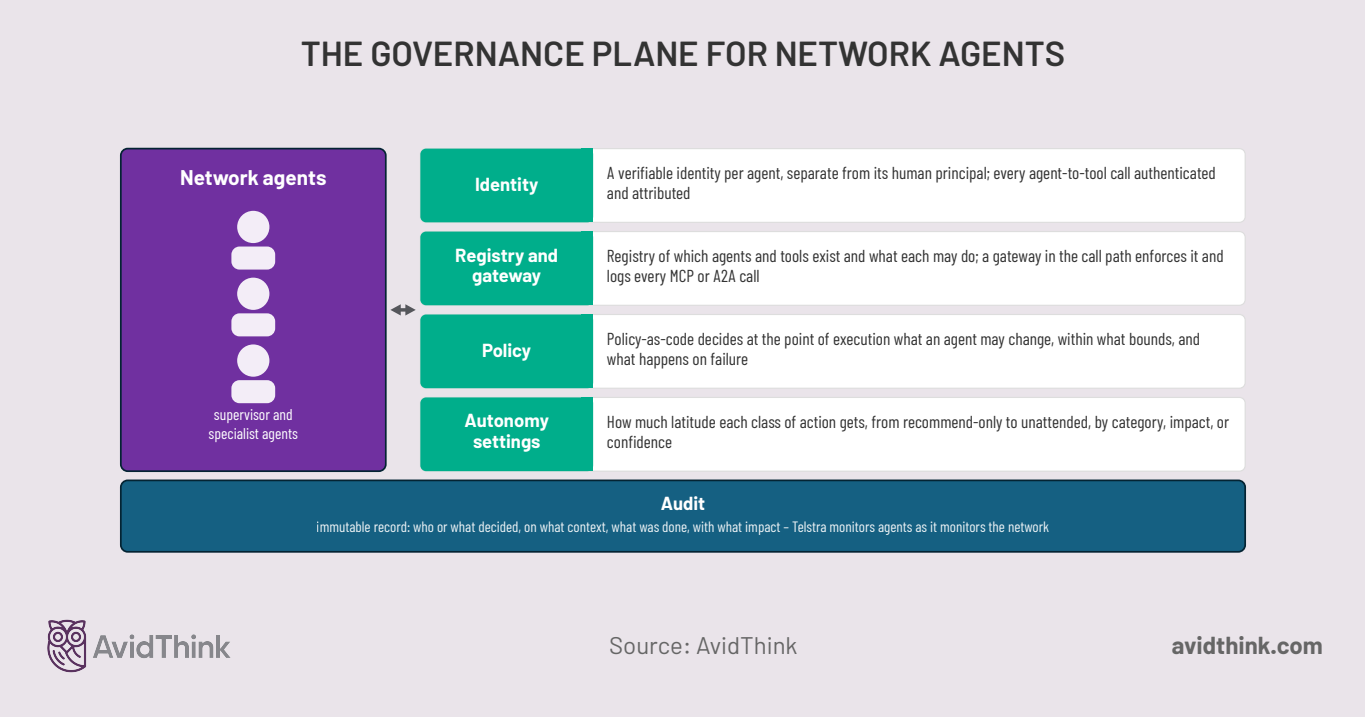
[63] K. Krogh Andersen (Telstra), "How Telstra Is Reinventing Connectivity for the AI Era," FutureNet Asia 2026, Singapore, Sep. 23, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[64] H. Kardme and N. Kumar (Cisco), "Cisco Secure Networking: Securing the Agentic Workforce," Networking Field Day 41, Oct. 9, 2026. https://techfieldday.com/video/cisco-secure-networking-securing-the-agentic-workforce/

©2026 AvidThink LLC. All Rights Reserved.

gateway with audited access rules in production,[65] and Itential's FlowMCP Gateway extends Itential's authentication and policy enforcement to other vendors' MCP tools.[66]

- **Policy.** Policy-as-code decides, at the point of execution, what an agent may change, within what bounds, and what happens on failure. Mobile platform vendor Rakuten Symphony argues for a policy-as-code "safety envelope" enforced at the point of execution.[67]
- **Autonomy settings.** Operators set how much latitude each class of action gets, usually by category and by impact (blast radius) or confidence threshold. Cisco's planned autonomy controls are set by action type, network domain, and change window,[68] and Amdocs lets operators set the autonomy level from an agent control tower and escalates uncertain cases to a human;[69] HPE's staged self-driving actions and Extreme's Coworker and Operator modes reflect the same idea.



Trust: Verification, Explainability, and Rollback

From Approval-Based to Evidence-Based Governance

Today's governance assumes a human in the loop: propose a change, wait for approval, then implement. As Blue Planet's Balaji Subramaniam argues in his interview later in this report, the approval-based approach won't scale with autonomous networks. The alternative he describes is a human-on-the-loop (HOTL) model built on collected evidence: an immutable record of who or

[65] J. Taaffe, "How telcos are building safe environments for AI agents," TM Forum Inform, Feb. 24, 2026. https://inform.tmforum.org/features-and-opinion/how-telcos-are-building-safe-environments-for-ai-agents

[66] "Itential Brings Governed AI Agents to Enterprise Infrastructure with FlowAI General Availability," Itential, Jun. 1, 2026. https://www.itential.com/news/itential-brings-governed-ai-agents-to-enterprise-infrastructure-with-flowai-general-availability/

[67] T. Ranjan (Rakuten Symphony), in "Rakuten Symphony on balancing AI agent autonomy with network control," TelecomTV, Jun. 16, 2026. https://www.telecomtv.com/content/apis/rakuten-symphony-on-balancing-ai-agent-autonomy-with-network-control-55675/

[68] J. Vaccaro, "From signal to action: The next step in Cisco AgenticOps," Cisco Blogs, Jun. 2, 2026. https://blogs.cisco.com/networking/from-signal-to-action-the-next-step-in-cisco-agenticops

[69] R. Porter (Amdocs), in "Amdocs on the journey to autonomous networks," TelecomTV, Sep. 23, 2026. https://www.telecomtv.com/content/the-ai-native-telco-forum/amdocs-on-the-journey-to-autonomous-networks-56324/

©2026 AvidThink LLC. All Rights Reserved.

what initiated a decision and why, what context informed it, what was done, and with what impact. He also flags the questions HOTL leaves open: what is the blast radius of the changes, and are there snapshots to roll back to (and how far back) when the human observer spots an issue later?

Rollback deserves more attention than it gets. A configuration rollback is the easy case: with intent in Git or a golden configuration, the previous state is merely a commit away. Examples here include Nokia EDA's dry run with rollback[54] and Telekom Malaysia's dedicated revert agent.[70] But a stateful rollback is harder: once BGP has reconverged, sessions have dropped, or traffic has drained to another site, restoring the configuration doesn't necessarily restore the network to where it was.

The practical answer so far is to classify every action class before we delegate it to agents. Categories include: *reversible* (within some acceptable threshold), *reversible with impact* (a configuration change whose reversion may cost a reconvergence or a maintenance window), and *irreversible* (major topology change, a drained site, a deleted object). For irreversible operations, regardless of the reasoning agent's confidence score, humans should be involved.

Nonetheless, we're seeing "auto" mode in many of today's coding agents, where the harness makes a judgment call about when to loop in the human for a decision. We expect the same "auto" modes to appear in network AIOps tools, and it will take a few iterations before we learn the best practices.

Building the Trust Ladder

This year, vendors (both telecom and enterprise) have converged on a trust ladder whose rungs are the five operating modes we defined earlier: (1) observe and explain; (2) recommend; (3) execute with individual approval; (4) execute approved action classes within limits; (5) operate a defined closed loop under policy and continuous verification. Scope is the width of each rung: a mode is earned first for one scenario, then for a domain, and only then across domains. The lower rungs map closely to today's coding agents' permission modes.

Most users are comfortable with Modes 1 to 4, and Mode 4 is where every networking GA product stops today. Only sophisticated organizations (mostly leading telcos and hyperscalers) have reached Mode 5, each with its own proprietary builds and each for a single scenario; nobody we know of has achieved Mode 5 across domains.

Trust will be earned over time: as positive examples accumulate, the industry will become comfortable raising the trust level for autonomous operations.

Mode 5 examples are few to date. The telco AIS closes 95% of alarms into jobs with no human intervention, with automatic rollback; it is operator-built, and it did not disclose how much of that loop is model-driven rather than rule-driven.[71] Deutsche Telekom's RAN Guardian, agentic and operator-built with Google, pre-checked 611 sites for the Carnival festival and adjusted five at peak.[12] Verizon's 70 million closed-loop configurations are often cited alongside these. Still, they are conventional closed-loop automation, cumulative across its platforms: evidence of a mature deterministic layer, not of agentic adoption.[37] Telekom Malaysia addresses the rollback half of the problem directly, with a dedicated revert agent for agent-initiated changes.[70] Each Mode 5 system is still limited to one scenario in one domain. Trust will be earned over time: as positive examples accumulate, the industry will become comfortable raising the trust level for autonomous operations.

[70] M. Tajul Mohamed Sultan (Telekom Malaysia), "The Role of Agentic AI in Autonomous Networks," FutureNet Asia 2026, Singapore, Sep. 22, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[71] W. Wattanasap (AIS), "AIS Case Study: From Network Automation to Agentic Intelligence," FutureNet Asia 2026, Singapore, Sep. 23, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

©2026 AvidThink LLC. All Rights Reserved.

Interoperability: MCP, A2A, A2A-T, and the Open Agentic Ecosystem

MCP: Evolving and Maturing

A year ago we called MCP "not a silver bullet and somewhat overhyped," flagged its missing security considerations, and said we expected the community to close the gaps quickly. It did. Anthropic donated MCP to the Linux Foundation's Agentic AI Foundation in December 2025.[72] The 2026-07-28 specification revision is a breaking redesign: stateless, handshake-free sessions; a multi-round-trip pattern replacing server-initiated requests; and the deprecation of several capabilities. Authorization stays on OAuth 2.1 with OAuth or OpenID Connect discovery.[73]

This year, adoption is no longer in doubt. Every major networking vendor ships MCP servers or gateways. The telecom operators we spoke with use it the same way, to integrate across systems.

An MCP server can abstract complexity and add value, or it can wrap APIs poorly and send agents into token-eating reasoning binges.

Nonetheless, MCP servers don't replace the need for consistent, well-specified APIs. An MCP server can expose useful services that abstract complexity, or it can merely wrap APIs without sufficient value-add and drive the agents consuming them into token-eating reasoning binges.

We still have data security and governance concerns, especially about early MCP implementations that haven't been updated recently, and expect more breaches or information leakage in 2026 and 2027 from misconfigured, poorly designed, and badly coded MCP services.

A2A and Its Telco Dialect

A2A reached version 1.0, its first stable specification, in March 2026 under the Linux Foundation,[74] with signed Agent Cards for cryptographic identity.[75] In networking, its uptake is slower than MCP's; we have not seen much discussion of it in an enterprise AIOps context.

The TM Forum has a telco-specific development called A2A-T: a telecom extension of A2A typed against telco vocabulary and frameworks.[76] Amdocs puts the savings in inter-agent tokens at 40–70%, a vendor estimate rather than a measured result.[77] Meanwhile, Telenor Sweden's architects are not convinced the industry needs a telco-specific MCP or A2A.[65] This debate will continue into 2027.

[72] "Donating the Model Context Protocol and establishing the Agentic AI Foundation," Anthropic, Dec. 9, 2025. https://www.anthropic.com/news/donating-the-model-context-protocol-and-establishing-of-the-agentic-ai-foundation

[73] "Key Changes," Model Context Protocol specification, revision 2026-07-28, Jul. 28, 2026. https://modelcontextprotocol.io/specification/2026-07-28/changelog

[74] "A2A Protocol Surpasses 150 Organizations, Lands in Major Cloud Platforms, and Sees Enterprise Production Use in First Year," The Linux Foundation, Apr. 9, 2026. https://www.linuxfoundation.org/press/a2a-protocol-surpasses-150-organizations-lands-in-major-cloud-platforms-and-sees-enterprise-production-use-in-first-year

[75] "A2A Protocol Ships v1.0: Production-Ready Standard for Agent-to-Agent Communication," A2A Protocol Community (Linux Foundation), Mar. 12, 2026. https://a2a-protocol.org/latest/blog/2026/03/12/a2a-protocol-ships-v10-production-ready-standard-for-agent-to-agent-communication/

[76] "Agent Fabric: A2A-T Runtime – Phase III," TM Forum Catalyst C26.0.910. https://www.tmforum.org/catalysts/projects/C26.0.910/agent-fabric-a2at-runtime-phase-iii

[77] J. C. Mendez (Amdocs), "Operationalizing the AI-Native Telco – Autonomy Without Blind Spots," FutureNet Asia 2026, Singapore, Sep. 22, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

©2026 AvidThink LLC. All Rights Reserved.

Runtimes, Hyperscalers, and Sovereignty

Increasingly, on the telco front, hyperscaler agent runtimes are becoming the delivery vehicle for telco autonomy stacks. Nokia is bringing its Autonomous Networks Fabric to AWS (availability expected late 2026)[78] and its assurance agents run on Google Cloud;[35] Blue Planet's Agents-as-a-Service runs on AWS with Amazon Bedrock underneath;[36] Ribbon's Acumen runs its agents on Bedrock AgentCore;[34] Amdocs pairs its Network Digital Twin and network-operations agents with AWS cloud and AI services;[79] Ericsson's agentic rApp-as-a-Service (rApps are RAN automation applications) is sold through AWS Marketplace;[80] DOCOMO is live on AgentCore;[13] Deutsche Telekom is on Gemini and Vertex.[12]

This arrangement runs in parallel with the sovereignty track that many of the same networking vendors are simultaneously pushing. Many telco vendors offer bring-your-own-LLM and self-hosted options, and large enterprise networking vendors are moving along the same path. The goal is to let enterprises and carriers alike exert some control over their assets, where processing happens, which model is used, and the governance mechanisms.

A Cross-Domain Perspective: Service Provider, Enterprise, and Data Center

As we take stock across all the major domains, we note the following common trends:

1. Architecturally, the agentic solutions across all the domains are converging on some form of the generalized architecture we discussed earlier.
2. Generally available vendor products in all three domains stop at pre-authorized action classes (Mode 4) for autonomy; only operator-built, single-scenario loops (Mode 5) in carrier environments go further; and no organization has fully autonomous cross-domain implementations.
3. Actions on the network today are bounded by a reliable deterministic execution layer that was built during the pre-GenAI, pre-agentic AI automation era.
4. MCP is a common tool surface used by every vendor and telco to pull in third-party data, or used along with APIs for external tool calling.
5. Establishing comprehensive and reliable sources of truth and context contributes more to outcomes today than pushing the boundaries in agentic technology, such as using the latest frontier reasoning model.
6. Governance has to be built before scaling, or is being rebuilt after sprawl (which we recommend avoiding).

[78] "Nokia, Amazon Web Services expand collaboration to deliver autonomous networks built for the AI era," Nokia, Jun. 24, 2026. https://www.nokia.com/newsroom/nokia-amazon-web-services-expand-collaboration-to-deliver-autonomous-networks-built-for-the-ai-era/

[79] "MWC 2026: Amdocs and AWS Announce Multi-Year Strategic Collaboration to Accelerate AI-Driven Telecom Modernization," Amdocs, Feb. 27, 2026. https://www.amdocs.com/press-release/mwc-2026-amdocs-and-aws-announce-multi-year-strategic-collaboration

[80] "Ericsson launches Agentic rApp as a Service on AWS to accelerate Autonomous Networks transformation," Ericsson, Feb. 17, 2026. https://www.ericsson.com/en/press-releases/2026/2/ericsson-launches-agentic-rapp-as-a-service-on-aws-to-accelerate-autonomous-networks-transformation

©2026 AvidThink LLC. All Rights Reserved.

What Differs Across Enterprises, Data Centers, and Carriers

The three domains differ in scale, technologies, and workloads. The table captures the key differences we observed.

	Service provider	Enterprise	Data center
Who orchestrates agents	Mix of operators and vendors: many tier-1s build their own orchestration; customers demand open interfaces	Usually the vendor platform (Cisco Cloud Control, Extreme Platform ONE, HPE Mist, Aruba Central), with agents bundled in subscriptions	GitOps and intent pipelines plus a vendor controller (EDA, HPE Apstra, Cisco Nexus Dashboard); coding agents as front end
Where changes come from	Planned updates, or response to tickets within an OSS system	Tickets from ITSM systems such as ServiceNow, alarms from monitoring systems	Data center or cloud-level orchestration that drives network tasks
How change is verified	Digital twin as design goal; SOP and GitOps process; few production twin KPIs	Pre-approved action classes inside change control; HITL otherwise	DevOps CI, lab replica, on-demand digital twin
Adoption evidence	Carriers publish outcomes; earlier claims over-optimistic; mostly single-domain	Vendor usage metrics; limited evidence	Very limited practitioner outage and MTTR stats
Sovereignty pressure	Strongest: processing locality, sovereign cloud	Regulated buyers: customer-managed keys, air gaps	Varies by market; on-site controllers valued for robustness

Barriers, the Human Element, and Recommendations

Barriers in 2026

Unsurprisingly, the barriers we first introduced in 2024 remain: trust, explainability, fragmentation, privacy, unclear ROI, and human friction. End users and vendors across domains are making progress because greater network autonomy can reduce costs while increasing robustness and agility. Meanwhile, here are major friction points we've seen recently:

1. **Governance is critical to establishing trust.** The attribution problem, unassigned accountability, and the reliability and alignment issues we discussed in previous sections reduce trust in agentic systems. The fall 2026 agent incidents don't help. Strong governance, both organizational policy and technology-enabled, needs to advance quickly.
2. **Data debt is still measured in years.** Telcos and enterprises around the world still describe an ongoing struggle to fix the underlying data layer. The lack of clean, appropriate data affects model performance and accuracy, and, with that, the reliability of reasoning and the effectiveness of the resulting actions.
3. **The business case needs to be established more rigorously as we move to agentic.** Automation will reduce operational expenses and improve reliability by reducing complexity. However, agentic systems are early, reasoning tokens can be expensive, and cleaning up the data layer costs money. We expect more proof points over the next year as agentic systems in specific domains start to generate measurable outcomes that translate into ROI.

©2026 AvidThink LLC. All Rights Reserved.

4. **Strength of executive support correlates with uptake.** Clear signs from leaders setting autonomy-level targets, as CxOs at global carriers have done, can accelerate progress. Unclear autonomous-operations goals make it harder for teams to push autonomy initiatives inside organizations.
5. **Sovereignty requirements gate the pace of adoption.** Data-residency laws and critical-infrastructure rules will push regulated buyers toward air gaps and customer-managed models.[81] However, the deployment of these private systems will be more bespoke and complicated than a multi-tenant cloud-based solution that benefits from economies of scale.

Two older barriers persist: a lack of skills, which the NAF survey ranks first (27%), and people-related barriers (organizational, cultural, and personality), which collectively poll at 44%.[27] Education and management buy-in are necessary to resolve these. Transitioning to AI-powered networking is a major change and requires change management techniques to achieve success.

The Human Element in AIOps

In 2025, we said the network engineer would become an AI operator who defines intent and manages agents. This is becoming reality, but the interface varies from coding harnesses in DevOps/NetOps workflows in data centers to web-based UIs for enterprise campuses to bespoke interfaces at carriers.

The network engineer is becoming an AI operator, but the interface of choice depends on the domain.

Meanwhile, incentives are starting to align with the new role: Telstra ties a metric called "AI value delivered" to the bonus of all 30,000 of its staff,[82] including its engineers.

Nevertheless, we advocate that organizations balance human engineers and AI-powered agents, delegating the best-fit tasks to each. Agents can help gather the necessary artifacts and perform first-level investigations. Experienced network engineers will have better judgment and instincts about where to look in specific failures. A judicious combination of the two, at least in the immediate future, likely provides the best returns to the organization.

Longer term, the open question is what the engineer's hand on the tiller looks like once agents outnumber people. Approving every action (HITL) will not scale, and watching every action (HOTL) will generate more than humans can digest. Governance will have to become good enough, and trust will have to build, to reach more autonomous networks without a human bottleneck, but with a human accountable for actions that cannot be easily undone.

Updated Recommendations by Domain

Our advice in 2026 includes making decisions around IP ownership in the context layer, establishing measurability for process improvements and ROI tracking, and ensuring sufficient governance to build trust. We start with recommendations that apply across all domains, then provide specific recommendations for enterprises, carriers, and data centers.

All Domains

1. **Own the context layer.** Especially for carriers and data center operators, where the context can provide competitive differentiation via lower cost of operations, understand what constitutes the ontology in your deployment. Fund the source of truth before expanding budgets on agents; measure timeliness as well as data and context cleanliness. Then decide who will own and maintain context (you or your vendor partners).

[81] H. Ölsner and C. Miene (NTT Germany / 50Hertz), "Workflow-Driven Automation Under Air Gap Compliance," AutoCon 5, Munich, Jun. 11, 2026. https://www.youtube.com/watch?v=zBvGy34XRUI

[82] G. Byrne (Telstra), "AI Agents in Action — Delivering Value Now," FutureNet Asia 2026, Singapore, Sep. 22, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

©2026 AvidThink LLC. All Rights Reserved.

2. **Insist on a deterministic execution layer.** Agents should trigger proven workflows, controllers, or a GitOps pipeline. This layer governs how an agent acts and reduces risk, but you will still have to check the plans that invoked the action for correctness.
3. **Govern agents like risky users.** Identity, registry, gateway, least privilege, per-action logs that identify the agent and its human principal (if one exists), and a validated undo path. This is more critical now given recent agentic misadventures.
4. **Make the twin or the lab a strong verification gate.** As the scale and impact of changes grow, you should mandate no network changes without a dry run. But remember that a passing digital twin or lab run shows feasibility; it is not a guarantee of a good production outcome. You still need post-change verification of the production network.
5. **Measure KPIs before and after.** Baseline MTTR, alert volume, operational costs, and token spend before the pilot and after.[83]
6. **Keep humans in the loop by design.** Understand the impact and reversibility of all tasks: classify every action class as reversible, reversible with impact, or irreversible before delegating it, and keep HITL on the irreversible class. Guard against rubber-stamping.

Enterprises

1. **Start with the platform you own.** Agents, context, and guardrails are arriving bundled with your incumbent vendor's platform. Turn on AI-assisted troubleshooting and root cause analysis first; that's where engineers report the clearest MTTR gains today. Then pursue automation actions starting from Day 2, then Day 1, then Day 0: diagnose, then recommend, then act. Agentic configuration and planning capabilities are improving but still early.
2. **Grant autonomy one action class at a time.** Allow autonomy only on narrow, reversible, well-understood classes of operations, such as wired-port remediation, and register it as a pre-approved change in your ITSM change control. Once you've built a track record, widen scope by action type, domain, and change window; keep HITL for everything else.
3. **Ask every vendor how it meters its AI.** Pricing ranges from "all you can eat" with caps, to inclusion in the platform subscription, to bring-your-own-LLM, where the token bill is yours. Know which one you're buying, what visibility you get into token use, and how that changes when agents move from answering questions to running investigations.
4. **Keep the source of truth yours, to avoid lock-in.** Your NetBox, Infrahub, Nautobot, or Git repository should remain the system of record that feeds the vendor's context layer. Ask how your topology, policies, and accumulated context can be exported if you switch vendors, and whether your data is used to improve shared models.
5. **Settle the NetOps–SecOps boundary before agents touch policy.** Decide in advance whose guardrails win when a NetOps agent wants to modify an ACL or reroute traffic to fix a performance issue, and encode that decision in policy rather than leaving it to escalation.
6. **If you're regulated, write sovereignty into the RFP.** Customer-managed keys, self-hosted or air-gapped deployment, and bring-your-own-LLM options are now available from the major vendors, but expect them to be more bespoke and to lag the cloud-hosted features.

Carriers

1. **Grade autonomy per scenario and rescore honestly.** Rescore yourselves on the TM Forum levels so your team's efforts match your actual process maturity. TM Forum's 20 vendor-agnostic Technical Solution Packages for Level 4 scenarios, published in October 2026, with 40 planned by mid-2027, are a good place to start.[84]

[83] BT International, "Resilient Networks in an Era of Disruption: Futureproofing Through Network Automation & AI," FutureNet Asia 2026, Singapore, Sep. 22, 2026. https://futurenetworld.net/events/futurenet-asia-2026/event-agenda/

[84] "TM Forum launches reusable blueprints to accelerate Level 4 autonomous networks," TM Forum (via EIN Presswire), Oct. 7, 2026. https://www.einnews.com/pr_news/947794903/tm-forum-launches-reusable-blueprints-to-accelerate-level-4-autonomous-networks

©2026 AvidThink LLC. All Rights Reserved.

2. **Treat the ontology as intellectual property.** Tier-1s can justify building their own and exposing it to agents through a governed interface such as MCP. If you buy from your OSS/BSS vendor instead, decide up front, contractually and technically, which parts of the ontology and the operational knowledge it accumulates remain yours.
3. **Budget tokens like headcount.** Set token allowances by role and use case, route the routine path to cheaper models, and save frontier models for hard tasks. Consumption is growing faster than per-token prices are falling.
4. **Centralize agent ownership before sprawl sets in.** Assign an accountable owner for the agent estate and treat agent behavior as an operational signal.
5. **Scale closed loops scenario by scenario, with rollback built in.** Pick event-driven scenarios with clear KPIs for high-autonomy operations. Repeat the pattern domain by domain first rather than reaching for cross-domain autonomy.
6. **Keep the runtime choice portable.** Hyperscaler agent runtimes are the fastest path to production, but sovereignty requirements will tighten (especially for carriers outside the United States and China). Insist on bring-your-own-LLM and self-hosted options, and keep the context layer and governance plane under your control so processing can move as regulations change.

Data Centers

1. **Put agents on top of the pipeline.** Your GitOps, CI, and twin or lab gate already form the deterministic execution layer. Let agents propose changes as intent in Git, and let the pipeline validate, verify, and deploy. No agent should have a bypass to the network elements.
2. **Fix the foundation before adding reasoning.** This year's strongest data center results came from deterministic automation, a good source of truth, and a digital twin. Get those right first; reasoning agents amplify a sound pipeline and a poor one alike.
3. **Extend AI from the build cycle into Day 2 operations.** AI is mostly used today to build and change the fabric. The next gains are in operations: correlating telemetry, explaining incidents, and recommending fixes in Modes 1 and 2 before granting execution rights.
4. **Measure the network by what the workloads need.** As AI workloads dominate, shift KPIs from uptime toward GPU utilization and the network's contribution to it (job completion time, network fabric congestion), and baseline both before introducing agents.
5. **Hold vendors to your operating model.** Ask your fabric controller vendor for declarative, intent-based APIs and MCP interfaces your harnesses can drive, and evaluate its upcoming operator modes against your existing CI gates.

Conclusion: The Path to 2027

A year ago, we wrote that yesterday's pipe dreams were becoming operational realities. And they now are, one scenario at a time. We still need to install and improve guardrails and rely on HITL for most processes and HOTL for some. Likewise, we are still reliant on (and should continue to rely on) a deterministic action layer. As the state of the art in agentic harnesses for coding and enterprise productivity improves, the networking industry will continue to borrow from it.

Over the next 12–18 months, we expect enterprise products to bring domain-level autonomy (campus Wi-Fi or another currently cloud-managed, well-understood domain). On the carrier side, we expect network-wide averages to improve but still stay below Level 3.

Even as we enter 2027, we expect humans to approve irreversible (or hard-to-reverse) operations with large impact radii. Likewise, the context layer, not the model, will decide which networking vendor solutions win, and token budgets will still worry CFOs. The industry can still reach full network autonomy in the 2029–2031 timeframe, within the five-to-seven-year horizon we set in 2024, but we'll revisit that projection next year, in our 2027 edition. In the meantime, reach us at research@avidthink.com.

©2026 AvidThink LLC. All Rights Reserved.

SPONSOR INTERVIEWS SECTION FOLLOWS

The following section contains interviews with sponsors who helped make this research possible. Sponsors do not have editorial control over report content.

©2026 AvidThink LLC. All Rights Reserved.

Q&A with Monica Kumar

EVP and Chief Marketing Officer, Extreme Networks



We have been using AI and algorithms to manage networks for some time. What is fundamentally different about agentic AI compared with traditional AIOps and generative AI?

There has been a generational shift. Traditional AIOps was about detecting a problem: when an anomaly occurred, it correlated events across the network and told you what the problem was, but that is where it stopped. Generative AI meant you no longer had to go through five dashboards to investigate; you could ask a question in plain English and be told what was wrong. Agentic AI takes this to a different level. You can not only understand what the problem is and why it happened, you get a recommendation about what to do next, and the AI can take action on your behalf within the guardrails your organization establishes. At Extreme, our network management portfolio, Extreme Platform ONE™, brings together network telemetry and operational information. Embedded within it is Extreme Agent ONE™, which has two modes: Agent ONE Coworker mode reasons about network conditions, explains why a problem is occurring, and recommends what to do next; Agent ONE Operator mode goes a step further and can autonomously take action.

If AI starts making changes to the network, how do you make sure it can be trusted?

We do not believe in giving AI unrestricted access to network operations. It is not simply about autonomy; it is about trusted autonomy. Agentic AI can reason based on the context you give it, but without the right context it cannot give accurate answers or take appropriate actions. Trust is the next step, and that is where people come in: they establish the guardrails within which the AI operates. The more deeply agentic AI is embedded in an organization, the more important those guardrails, safety controls, and governance mechanisms become. Take our customer, New Perspective Senior Living, which operates about 48 senior living communities across eight states and plans to add another 100. It is a highly digitized healthcare environment: patient care, nurse-call systems, cameras, keypads, access logs, and electronic medical records all depend on the network. To bring communities online quickly, they created preconfigured, AI-based templates for deploying networks with Extreme Platform ONE, and launched nine new communities in a week, versus months previously. Importantly, the guardrails came from the organization: they defined what the AI should do and the parameters within which it could operate. That is how you establish trust.

Why does Extreme Platform ONE matter? Couldn't an AI agent simply connect to existing management tools and data sources?

Technically it could. But the question is whether it operates with the right context for your environment and within the appropriate guardrails. Problems can occur in many places: Wi-Fi, an access point or switch, a security policy, authentication, the application itself, even the underlying hardware. Extreme Platform ONE brings together telemetry across devices, users, applications, and infrastructure and makes it available to the agent, so it has the data and context it needs to reason about what is happening. It is not about having multiple agents running around. Unless they have the right context, the right data, and appropriate guardrails, they will not deliver the outcomes organizations need. The more agents you use, the more important the underlying platform, data, and context become.

Does agentic AI simply make troubleshooting faster, or does it change how networks are operated?

It fundamentally changes how networks operate, from reactive to proactive. Traditionally something breaks, somebody notices or gets a notification, and an investigation across multiple dashboards begins that can take hours or days. With agentic AI and the right platform context, the network itself recognizes that something is wrong. Agent ONE Coworker mode can proactively notify you through its Nudge capability, investigate, and come back with recommendations; with Agent ONE Operator mode, it can take corrective action if that is what you want. We are not taking the old workflows and making them slightly faster; we are shifting network management from reactive to proactive operations.

Nearly every networking vendor is introducing AI agents. What is different about Extreme's approach?

Anyone can have an AI agent. The important question is what that agent is connected to, and what context it uses to make decisions. Our agents are grounded in the operational data and network context available through the platform. And the more autonomy you give your AI, the more important governance and safety controls become. It's not who has an agent, but who has an agent that can operate within the context and safety controls of your environment. That makes it a platform question, not a chatbot question.

How will agentic AI change the role of network professionals?

Agentic networking is not about reducing the number of people operating a network; it is about changing how they work and where they spend their time. Networks are becoming more complex, yet headcount and budgets are flat. With Agent ONE Coworker and Operator, teams can reclaim hours or even days otherwise spent investigating issues and looking through dashboards. Rather than firefighting, they can focus on work that is more strategic for their organizations.

gluware®

Q&A with Ernest Lefner

On automation-resistant networks, Proof of Agentic Trust, and turning CVEs into action



Ernest Lefner
Chief Product Officer, Gluware

Network automation has been on the agenda for years. Why has it suddenly become urgent?

For many years, CTOs and infrastructure managers treated network automation as an unsolvable problem that sat off to the side. With Anthropic's Mythos and advanced frontier models finding vulnerabilities faster, it is now a center-of-the-table problem, and the risk to infrastructure is visible all the way to the top of the house. You have to fight fire with fire. The answer is a platform-based approach that stops focusing on the minutiae of automation, starts focusing on business outcomes, and uses AI to get there.

Most networks aren't clean builds. What makes them "automation resistant," and why not just start over?

Different parts of a business have different network requirements, so up to 30% of a network's configuration is bespoke, and snowflakes are everywhere. You need a technology stack that can data-model the network and melt those snowflakes so automation can come through. Starting over is expensive. Many analysts now advise stretching depreciation from five years to seven or eight, partly because memory has become so costly. And if you build greenfield without solving automation, drift sets in immediately and the new environment wilts. You need this capability whether you buy new or keep what you have.

Letting AI change a production network takes a lot of trust. How does Gluware establish it?

Experimenting, such as generating code or pulling examples from the internet, is a fine way to get introduced to AI. But a high-risk enterprise environment needs to take a step back. Our Proof of Agentic Trust framework starts from what we already know: your standards and how every device is actually configured. When we make a recommendation, it reflects both best practices and your network as it is today, snowflakes included. Take a large EVPN-VXLAN rollout on the core. The AI has to understand how every component must be configured to your specifications, so we build the rules engine and tool calls that apply the right checks and validations. One agent proposes a configuration; a separate agent, with no bias toward what was created, validates it independently and queues it for review in the governance process. Think of LiDAR in a self-driving car: we can see what's out there. Good certification principles from the past 30 years still apply.

Where should teams start?

Crawl, walk, run, starting with adoption. Bend the branch lower by tackling what makes automation hard. For example, AI can build the audits that confirm your network is in a state that's safe for change, validated and tested with proof outputs, instead of users writing them by hand. The low-hanging fruit is stabilizing your environment: knowing where you stand against your standards and what your lifecycle management looks like. Then you package known-good audits, configurations, and state assessments and run them quickly and repeatably.

How does Gluware's Titan Exposure Management turn vulnerability data into action?

It's an AI product built on Gluware Titan AI that looks across your inventory and pulls in multiple external feeds. When a CVE is published, it reads the advisory, runs the appropriate show commands, and checks your whole fleet for vulnerable configurations. You might have a thousand switches of one model, but only a hundred running the affected feature. That turns a monumental problem into a manageable one, in hours rather than the weeks a manual investigation takes. Where the advisory includes recommended configurations, it recommends compensating controls you can push immediately, or you can schedule upgrades. Gluware is API-rich and can run almost headless, so results flow to ServiceNow, NetBox, Qualys, and other systems through MCP or data syncs such as CMDB sync. A full IDE also lets customers build their own applications on the platform.

Why Gluware, and what should network leaders do now?

Gluware represents more than a million development hours. DIAL, our Device Interaction and Automation Layer, handles two-way communication with devices, so you have one consistent way to build and deploy automation at scale. With self-built Ansible and Python, ten platforms running ten operating systems could mean a hundred scripts for every task. With Gluware, you focus on the intent and the outcome. Everything is a link in the chain: to upgrade faster, devices must be correctly configured, with failover working and standards met. So we take the problem end to end: stabilize, make change safe and predictable, understand the lifecycle, and then fire off upgrades or compensating controls, rather than treating each problem as a separate task. The industry has changed irreparably, and we are in a buffer zone where you can rethink how you deliver services. Our goal is a platform that works with your equipment regardless of vendor, provides a uniform experience across the enterprise, and brings in AI that protects the network.

Many of our customers have leveraged Level 1 and Level 2 capabilities delivered by Nokia NSP, and now we are helping them move toward Levels 3 and 4.



Q&A with Mike Thompson,

VP, Head of Network Automation Product Management

Agentic AI has become the dominant theme in network operations this year. What does an AI agent actually mean in this context, and why is now the moment?

Today's AI agent is autonomous: we give it a goal, and the agent assesses its environment, evaluates, plans, makes decisions, acts, and keeps iterating. In networking, the primary value comes from orchestrating diverse agents to solve different operational problems, securely and in a way that builds operator trust. Nokia is committed to helping customers through the journey to agentic operations. We recently added an agentic AI framework to Network Services Platform (NSP), our multivendor, multidomain, network automation platform. Operators can now deploy AI agents grounded in a real-time, accurate network view—enabling them to reason on trusted data and take guided, explainable actions within operator-defined policies and security boundaries.

Where are operators today on the TM Forum autonomy scale, and where do they expect AI to pay off?

TM Forum's framework is a useful map. Level 2, partial autonomy, is where the industry is centered today. At Levels 3 and 4, AI plays a much bigger role, ultimately leading to fully autonomous networks at Level 5. On the payoff side, analyst firm Appledore Research spoke with operators earlier this year, and they overwhelmingly pointed to reducing mean time to repair, optimizing capacity, improving productivity, speeding service activation, and reducing truck rolls. Many of our customers have leveraged the Level 1 and Level 2 capabilities delivered by Nokia NSP, and now we are helping them move toward Levels 3 and 4.

Many vendors claim AI-driven operations. What has to be in place underneath for agents to be useful?

AI without the right data isn't very valuable. Fortunately, NSP runs in many of the world's top service provider networks, so we have deep data on topology, inventory, service levels, and performance. We unify that data through ETL, curate, and structure it through an ontology. On top sits an agent and context framework that handles agent lifecycle, memory, context, role-based access control, guard rails, and validation. We are also model-flexible, we let customers pick their models. Token cost is increasingly a real concern. Running agents in a lab is one thing; running them at scale is another. Customers want the right results at the right cost and a real return on investment—we can help deliver that.

What agents and interfaces are Nokia bringing into NSP?

We are building purpose-built agent runtimes, starting with troubleshooting agents and control-plane assurance agents, and we will keep adding specialized agents. These agents draw on analytics from our data lake, such as predictions, baselines, and anomaly detection, and tie into automation flows for closed-loop operation. Copilots let operators ask questions in natural language. In addition, our new Agentic Command Center brings natural-language querying into the user interface, so operators can modify maps, retrieve information, and get help with troubleshooting and insights. For interoperability, NSP has an MCP interface; we are working on A2A (agent-to-agent), and we continue to support Kafka and other established APIs.

Operators rarely run single-vendor, single-domain networks. How and where does NSP fit?

NSP is highly integrated but open. It spans infrastructure and service fulfillment, assurance and analytics, and deep control and optimization through our path computation engine. Its automation and customization capabilities let it operate in networks that include Cisco, Juniper, Huawei, and other vendors across both the IP and optical layers; with Infinera, Nokia is now one of the industry's largest optical vendors. For cross-domain autonomy, it works alongside Nokia's Corteca and Altiplano in fixed networks, MantaRay in mobile, WaveSuite in optical, and EDA and Digital Operations Center in core and data center.

What will it take to move operators to Levels 3 and 4, and what is your message to network engineers?

It is not simply about technology. It's about achieving outcomes. At Nokia, we work through use cases until they are repeatable, so our customers can implement them faster and with lower risk. For AI-forward customers, NSP is also a highly programmable, developer-focused platform that integrates well with tools such as Cursor (or other code harnesses) via MCP and APIs, supporting intent-based development. We want to empower our customers' engineers to combine their networking, AI, and programming skills. Ultimately, our goal is to help operators deliver better service faster, save money, and improve productivity. AI can bring tremendous value, but it is also complex, and operators need a partner with proven multi-vendor, multi-domain experience like Nokia.

Q&A





Ram Ramanathan
Vice President, Product Management
Ribbon

Chris McKenna
Global Head of Business Development – 5G Networks
AWS

What's driving service providers to rethink their approach to managing the complexity of network operations?

Ram: Networks are getting more complex with 5G data and voice rollouts and AI workloads from the edge to the core. At the same time, margins and ARPU are being compressed, and experienced experts are retiring. Networks are still run in silos, with each domain and layer handled separately and multivendor islands of context. If implemented the right way, AI can let carriers manage that complexity efficiently while improving margins and reducing churn. That's why Ribbon launched Acumen, our agentic AIOps platform.

Chris: We see the same signals. Networks have become too distributed and complex for people to manage at today's speed. Meanwhile, AI has become capable of reasoning, not just pattern matching. It can examine thousands of alarms from 15 or 20 different systems and make sense of them in concert. The shift is from detecting problems to explaining and resolving them.

Where should operators start in their adoption of AI operations?

Chris: With low-risk use cases. Nobody dives into fully automated closed-loop operations on day one. We see customers using AI effectively for lifecycle management and large, complex rollouts. Today that work is human-led with AI assistance; over time it will become AI-led with human oversight.

Ram: Agreed. The first is assisting NOC teams as these new technologies add complexity. The second is change management. Many network changes are still made manually, often on a leap of faith. A controlled, low-risk way to make changes is a natural early win.

What does Ribbon and AWS each bring to the solution?

Ram: Acumen is a platform on which many agents can be built, designed for value- and ROI-driven outcomes rather than AI for its own sake. Customers can create their own agents, and we tailor agents to each use case, because a NOC agent or change agent differs from one operator to the next. The architecture layers data, task agents, and orchestration agents, built on more than 25 years of Ribbon's knowledge and experience in the network domain.

Chris: The two are complementary halves. Ribbon brings the domain intelligence: topology, protocol knowledge, and what a healthy network looks like. AWS brings AI tools that reason over that knowledge at scale. For example, the network can be represented as a live digital twin, a graph in Amazon Neptune, so AI sees relationships rather than isolated alarms. Amazon Bedrock provides the right model for each agent's task.

Walk us through an example. Say the carrier sees service degradation but no alarm triggered. How can the system find the cause?

Ram: Today's tools are siloed. A mobile carrier has separate tools for RAN, packet core, and IMS, with limited correlation between them. Acumen correlates operational data across Layers 0–7 and across domains, combining telemetry from network devices with topology, geography, and traffic patterns. It baselines the network first, then uses what it learned, plus the operator's knowledge bases and ticketing data, to reason about anomalies. And we are building our platform to do this in real time at scale, rather than in a post-mortem.

Chris: An operator queries the system in plain language, which kicks off a team of agents on Amazon Bedrock AgentCore. Specialist agents for RAN, transport, packet core, and voice core run in parallel, pulling telemetry, walking the topology graph, and retrieving runbooks and past incidents. A reasoning agent synthesizes a single root cause and formulates the explanation. The Ribbon platform then executes the domain-specific remediation. Work that used to take multiple people hours or even days reduces to minutes.

When does the system move from diagnosing problems to fixing them on its own?

Chris: That's the million-dollar question. Diagnosing a live network and changing it are very different things. We design for trust in stages: guardrails in Bedrock, with Ribbon adding its own on top; human approval by default; autonomy granted per action type. Full observability is critical, because agents need to see everything, and every step must be logged with its evidence. Looking two or three years out, I think we'll start to see the first instances of closed-loop remediation.

Ram: It's crawl, walk, run, with humans in control for the next couple of years through copilot-like projects such as a NOC copilot and a change copilot. Platforms such as Acumen need policies defining what network data, memory, and context that agents can access, and for closed-loop feedback. These controls add to carriers' existing role-based access, NOC procedures, audit trails, and compliance rather than replacing them. Trust has to be built in, not added as an afterthought.

If we spoke again in two years, what would have changed?

Ram: My best guess is a steady move toward intent-based automation. Someone asks for a transport link between two points with certain SLAs, and agents make it happen, then perform an audit to ensure that the SLAs are met. Humans set business priorities; agents do the work.

Chris: Routine diagnosis and remediation will become autonomous and predictive, and operators will move into supervisory roles that handle novel scenarios. The milestone to watch is an operator trusting closed-loop remediation to run unattended in production.

blueplanet
a division of Ciena

In conversation with Balaji Subramaniam

Why autonomous networks need evidence-based governance



Balaji Subramaniam
Director, Product Management
Blue Planet, a division of Ciena

Blue Planet, a division of Ciena, provides multivendor inventory, orchestration, assurance, and configuration and change management software, with AI embedded directly into OSS workflows to enable more intelligent, trusted, and autonomous network operations.

Autonomous networks come up in nearly every service provider conversation today. What is the issue you hear about most?

Governance. We talk with CSPs around the world, and governance keeps coming up in executive and boardroom conversations about autonomous networks. Operators are trying to understand how to maintain visibility, control, and trust as they introduce more AI-driven automation. At Blue Planet, we see that journey starting with trusted network and service data, combined with orchestration and configuration and change management across the service lifecycle. Together, these capabilities give operators a more practical path to autonomy without giving up control.

Why can't today's governance processes simply carry over to AI-driven operations?

Today's governance model is built around a human in the loop. You make a change, wait for approval, and then implement it in the network. That approval-based approach won't scale with autonomous networks, given the volume and complexity of AI-driven decisions and their potential impact. Governance needs to shift toward a human-on-the-loop model, moving from approval-based to evidence-based governance. Operators will need to rethink where they place guardrails and checkpoints, embedding governance into automation rather than treating it as a separate process.

What is the biggest governance risk when AI is taking action in the network?

Most people ask whether AI is making the right decision. That is not enough. You need to understand the consequences of what AI is going to do. It isn't making isolated decisions; it's making sequences of decisions that build on each other and can impact the network at scale. So what's the blast radius if something goes wrong? How far do you roll back? Do you undo the last action, or roll forward and correct things from that point? And because much of the context AI relies on is ephemeral, you need to capture that context so you can reconstruct the decision and verify that AI started from the right place.

What does "evidence" actually mean in practice?

Today, after an incident, teams dig through logs, piece things together, and eventually determine what triggered the event. That won't work at the speed and scale of AI-driven operations. Operators need immutable evidence of every decision: who or what initiated it and why, what network context and parameters informed it, which actions were taken, and what impact they had on the network. That evidence must be captured in a persistent, auditable record that cannot be altered.

How is Blue Planet building this into its products?

We're building these capabilities into Blue Planet Configuration and Change Management, which integrates with our broader portfolio. The goal is to capture not just what changed, but who or what initiated it, why it happened, and what impact it had, creating an immutable, auditable record of network changes. We're also introducing pre-built AI agents for drift detection, compliance validation, and change risk assessment. By connecting these capabilities with inventory, orchestration, and assurance, we can bring greater visibility, control, and traceability to AI-driven operations.

What is the takeaway for operators planning their autonomy roadmap?

Governance becomes even more essential when operations no longer require someone to click "approve." It has to be embedded in how autonomous systems make decisions, execute changes, and provide evidence of their actions. Governance cannot be an afterthought. Autonomous networks require trusted data, intelligent automation, and governance working together. Robust configuration and change management gives operators the visibility and control to move toward autonomy progressively and responsibly.

Appendix: Vendor Profiles

This appendix profiles select networking vendors that we've engaged with, or that have surfaced in our conversations with end users. They are grouped by functional role—enterprise and data center infrastructure platforms; automation, orchestration, and source-of-truth specialists; observability and AIOps specialists; telco network infrastructure, OSS/BSS, and RAN ecosystem; and horizontal AI and workflow platforms. Within each category, they are listed alphabetically. For each, we give a description and the most significant developments since our last edition. This section includes vendor-specific detail; the main report references vendors only when they illustrate architectural points. The automation, observability, and source-of-truth specialists appear this year because of their role in the ongoing architectural evolution. You should verify the information in this section directly before making any purchase or shortlist decisions, given the rapid advances in this field.

APPENDIX | A. Enterprise and Data Center Infrastructure Platforms

Arista Networks

Arista positions itself as a client-to-cloud networking provider across AI and enterprise data centers, campus and routing, built on EOS and the CloudVision management plane, and markets AVA as a multi-domain agentic AIOps platform. The major update since our 2025 profile, which covered AVA as a chatbot-plus-insights layer over the NetDL data lake, is that AVA was recast in December 2025 as a unified agentic framework: NetDL as the cross-vendor data lake, AVA as the agent layer, and MCP as the surface for third-party agents. AVA autonomously investigates from curated "skill documents," but every corrective change requires administrator approval under role-based access control (RBAC) and change control. A year after the VeloCloud acquisition closed, July's Edge Threat Management release put Ask AVA inside VeloCloud Orchestrator as a policy explainer and pre-change traffic simulation. However, CloudVision and VeloCloud remain separate management planes. Arista monetizes AVA through CloudVision subscriptions (about 3,000 cumulative customers).

Cisco

Cisco positions itself as critical infrastructure for the AI era, and its operations message is now AgenticOps. A year after we profiled AgenticOps at launch, the major update is a broad platform with a regular release cadence. AI Assistant for Networking reached general availability in October 2025; Deep Reasoning investigations followed, at roughly 150,000 queries a month by Cisco's count. In June this year, Cisco unveiled Cisco Cloud Control, a single login and data layer for humans and agents across Meraki, Catalyst, Nexus, Splunk, and Webex, included with most platform licenses, alongside AI Canvas. AI Canvas provides an Agent Builder that integrates with 50+ platforms natively or via MCP, and provides a per-category "autonomy dial." Agentic Actions is in beta, and an Agent Digital Twin is in alpha. Cisco reported more than 20,000 active AgenticOps users in six months, and Cisco IT (which uses its own products in-house) says its deployment cut mean time to resolution from seven days to two. All Agentic Actions changes require approval today; GA of Agentic Actions and the Agent Digital Twin are the next milestones.

©2026 AvidThink LLC. All Rights Reserved.

Extreme Networks

Extreme Networks has been one of the most ambitious vendors in its AI-for-networking roadmap, rapidly evolving from conversational AI to a proactive, context-aware agentic architecture backed by growing customer adoption. Introduced in May 2026, Extreme Agent ONE builds on Extreme Platform ONE's unified data foundation, combining multi-model AI orchestration, a networking knowledge graph, and specialized agent skills. The company leverages hyperscaler frontier models in its stack, allowing it to swap between models based on performance over time. Its knowledge graph serves as a key differentiator, grounding AI reasoning in network context and domain expertise, while Extreme Exchange supports its governance and extensibility strategy. Extreme Agent ONE operates in two modes: Extreme Agent ONE Coworker mode, which proactively investigates issues, delivers Nudges, and recommends actions with human oversight; and Extreme Agent ONE Operator mode, expected in the first half of 2027, which is designed to autonomously execute authorized tasks within defined governance boundaries.

Extreme Agent ONE Coworker mode became generally available in September 2026, included with Extreme Platform ONE subscriptions. Its capabilities include proactive Nudges, AI-assisted troubleshooting, radio resource management insights, conversational access to network data and knowledge, and streamlined support escalation. Extreme reports up to 15× faster troubleshooting, up to 50% shorter onboarding, and an estimated 1,000 troubleshooting hours saved by customers in its first month. Extreme Platform ONE exceeded 30% of subscription bookings in its first year, with Q4 bookings doubling quarter over quarter, and passed 50% in fiscal Q1 2027; Extreme Fabric continues to anchor larger enterprise wins. Extreme further detailed its agentic AI architecture and governance strategy at its AI Summit in Amsterdam on 20 October 2026.

HPE Aruba Networking / Juniper

HPE Networking, which absorbed Juniper Networks in 2025, markets "the self-driving network" across HPE Aruba Networking (Central, CX switching, access points, and SASE) and HPE Juniper Networking (Mist AI, Marvis, Apstra, routing, switching, and security) for campus, branch and data-center customers such as Mercedes-AMG PETRONAS F1 and Norfolk Southern. The major update since our 2025 profile is a single AIOps roadmap across HPE Aruba Central and Mist ("build once, deploy twice"), with one AIOps team driving both experiences. Marvis Actions have arrived on Aruba Central (starting with wired-port remediation), and Aruba CX switches now have Mist management. HPE has introduced a five-stage autonomy model, with the highest stage acting only with explicit permission, targeting Marvis Actions across wired, wireless, and SD-WAN in Central by end-2026. Current AI actions include RF capacity optimization and rogue-DHCP remediation. Juniper's Apstra is now HPE Networking Data Center Director with agentic root-cause analysis on the Marvis engine.

©2026 AvidThink LLC. All Rights Reserved.

APPENDIX | B. Automation, Orchestration, and Source-of-Truth Specialists

Aviz Networks

Aviz Networks, a San Jose startup backed by Cisco, NVIDIA, Celestica, and Keysight, sells vendor-agnostic SONiC software plus an AI-NOC and Network Copilot to enterprises, neoclouds, and telcos that want to "network like hyperscalers." Noted last year for Network Copilot's private-AI positioning, Aviz rebuilt it as a multi-agent AI NOC: at Networking Field Day 40, a super agent coordinated sub-agents running root-cause analysis across SONiC, Arista, and Cisco devices, pausing for a human when sub-agents disagreed, with critical changes requiring approval. Customers can deploy fully on-premises with open-weights models.

BE Networks

BE Networks (Richardson, Texas) is an orchestration software vendor whose Verity platform provides vendor-agnostic, intent-based orchestration of multi-vendor SONiC and Cumulus IP-Clos fabrics for data-center, campus, and AI-infrastructure operators, with Hilton, KDDI, NTT DOCOMO, and the US DoD as references. Its 2026 triad adds Satori (observability) and SensAI, a "Digital System Engineer" that drafts change plans and executes them after human approval. Verity 6.6.1 shipped a native MCP server with RBAC. IREN and NVIDIA announced a DSX Air digital twin plus BE's Verity to pre-validate a 50,000+ GPU build.

Gluware

Gluware is a network automation and agentic operations software company that handles complex, multi-vendor enterprise networks. Customers include Accenture, Merck, Pfizer, Mastercard, and the US Social Security Administration. Titan, the company's flagship platform, provides a full-stack network automation solution including OS lifecycle management, configuration intent modeling, compliance and drift auditing and remediation, cross-platform orchestration, and agentic AI capabilities. Automation and AI features rely on Gluware's DIAL (Device Interaction and Automation Layer), which interprets and validates proposed agentic actions. Positioned as a key differentiator, DIAL provides semantic translation across 55+ network operating systems from 22 vendors and has been tested to handle 450,000 nodes. Titan's Proof of Agentic Trust framework, which pairs agentic builder and validation skills with deterministic execution, is matched by a commercial model that lets customers bring their own LLM with no token metering beyond the per-device fee. To support enterprise IT environments, Gluware integrates with popular networking inventory and ITSM solutions such as NetBox and ServiceNow.

Gluware recently expanded its solution portfolio with Titan Exposure Management, a new capability unveiled in May 2026. Exposure Management is an agentic network vulnerability tool that scores CVEs against network devices based on OS version, configuration, and operational state within a customer network and helps fix affected devices. It automates the transition from vulnerability detection to deploying fixes, whether through automated software patches, configuration changes, or other compensating controls.

©2026 AvidThink LLC. All Rights Reserved.

Hedgehog

Hedgehog, a Seattle-based startup, sells open-source AI network software for GPU clusters to neoclouds. Its Kubernetes-native controller continuously reconciles SONiC-based leaf–spine fabrics on white-box switches. Its solution drafts and manages resource definitions, which a deterministic reconciler then runs to compile intent into switch configuration. Hedgehog ships no agent and no MCP server by design. Its enterprise customer, Zipline, runs a 60+ server training cloud on it with no NetOps staff.

Inmanta

Inmanta, a Belgian KU Leuven spin-off whose intent-based service orchestrator is used by BICS, Proximus Global, Telenet, and Citymesh, positions itself as the deterministic execution layer under AI rather than as an AI vendor. Inmanta Service Orchestrator 9, GA in May 2026, added report-only resources, a compliance API, and an Open Policy Agent policy engine, a concrete mechanism for staged trust. For Inmanta customers, an external AI solution proposes while Inmanta executes within intent boundaries.

Itential

Itential, an Atlanta-based orchestration vendor serving Lumen, Colt, Deutsche Telekom, financial-services firms, and US federal agencies, positions itself as the agentic operations platform for infrastructure teams. It is a clear shipping example of bounded actuation. Itential's FlowAI, generally available since July 2026, adds FlowAgents that reason and plan but act by launching governed and proven Itential workflows. FlowAI includes a FlowMCP Gateway that extends Itential's RBAC, audit, and approval gates to third-party MCP tools such as NetBox and Selector. Its carrier customer, Lumen, runs more than 350 live workflows and uses FlowAI in production.

NetBox Labs

NetBox Labs, the New York-based commercial steward of the open-source and widely deployed NetBox source of truth (10,000-plus organizations; more than $55 million raised, IBM and Salesforce among its investors), now sells an "infrastructure intelligence platform" for people and AI agents alike, with Arm, CoreWeave, and JPMorgan Chase as named customers. In June 2026, it launched NetBox Assurance (drift detection) and a Platform MCP Server with per-user authentication, branching, and approval gates. At its first NetBox Evolve conference in October, the MCP server reached GA, and six NetBox Agents entered preview with per-action allow, approve, or block settings. NetBox became a native source of record for NVIDIA's DSX OS configuration manager.

©2026 AvidThink LLC. All Rights Reserved.

Network to Code

Network to Code, the commercial steward of the open-source Nautobot source of truth, serves enterprises such as Chevron, Cox, Ørsted, and Intel and positions "data-driven network automation," rather than agents, as the key to its success. Noted last year in our report for an AI-powered automation suite, it spent this year becoming a product company. Nautobot 3.1 (April 2026) brought OS Upgrades and Operational Compliance apps, and recurring revenue grew 52% in 2025. NautobotGPT and a Nautobot MCP Server remain developer-facing capabilities.

OpsMill

OpsMill, the French company behind Infrahub, describes its product as "the AI-ready data management platform" for infrastructure automation, arguing that the platform matters more than AI model performance. Infrahub is a schema-driven, version-controlled graph source of truth; a $14M Series A in May 2026 cited TikTok in production and Eurofiber cutting service deployment from five days to fifteen minutes. Its MCP server implements guardrails as Git-style branching. Every agent write on Infrahub happens on a branch that a human must verify and then merge (HITL). The company claims over 600 deployments, including open-source installs.

APPENDIX | C. Observability and AIOps Specialists

Kentik

Kentik, the San Francisco-based "network intelligence platform" for flow, routing, synthetic, and cloud telemetry, counts ServiceNow, Zoom, and Booking.com among its customers. Its AI Advisor (November 2025) plans and executes investigations over a claimed trillion telemetry points a day, with Equinix identified as an early-access user. In March 2026, Kentik shipped a hosted MCP server exposing three conversational tools (`ask_question`, `ask_followup`, `get_session_history`), wrapping its agent rather than raw data, an agent-as-a-tool pattern.

Selector AI

Selector AI, a Santa Clara observability and AIOps vendor whose customers include Bell, Comcast, Lumen, and Singtel, now leads with Selector Foundry, an agentic NetOps platform. Noted last year as Itential's AI partner, it raised $32M at a $375M valuation in February 2026 (with claimed annual recurring revenue up 2.3× year-over-year). Its solution includes a managed MCP server (OAuth 2, RBAC) that exposes insights, telemetry, and workflows to external agents. In contrast, Selector Foundry's agents act on Selector objects such as tickets and dashboards under HITL. Closed-loop remediation remains on the roadmap.

©2026 AvidThink LLC. All Rights Reserved.

APPENDIX | D. Telco Network Infrastructure, OSS/BSS, and RAN Ecosystem

Aira Technologies

Aira Technologies, a Saratoga, California, AI software company backed by AT&T, Snowflake Ventures, and Neotribe, targets mobile RAN operations with its platform, Naavik One, pitched as an autonomous network driver for telecom. Naavik breaks down an engineer's troubleshooting expertise into small agents ("intuitions") that feed a reasoning agent for root-cause analysis. Naavik AppGen generates RAN rApps from natural-language intent, demonstrated with AT&T and Ericsson at MWC 2026.

Amdocs

Amdocs is a software and services provider to communications and media companies, with customers such as AT&T, PLDT, Telefónica, Telia, and Globe. In 2026 it introduced aOS, an agentic operating system for telecom spanning Amdocs and third-party BSS, OSS, and network environments. For Autonomous Network Operations, it pairs its Cognitive Core and telco-specialized agents with an Operational Digital Twin and Agentic Network Workflows for monitoring, detection, diagnosis, resolution, preventive operations, and change/configuration management. It works with several hyperscalers and integrates with customers' AI ecosystems. At PLDT, it reports 37% faster issue resolution and 94% fewer high-severity incidents.

Blue Planet (a division of Ciena)

Blue Planet, Ciena's OSS software division, serves communications service providers such as Telefónica Germany, Swisscom, Lumen, BT Group, and TDC NET with inventory, orchestration, assurance, and route-optimization products. Blue Planet frames its approach in three layers: a context layer built on a converged ontology spanning inventory, orchestration, and assurance, enriched with a "decisioning graph" that captures how operators have resolved past situations; reasoning agents embedded in each product; and supervisory "deep" agents that coordinate work across applications. Support for third-party models, agents, and open approaches such as MCP and A2A gives operators flexibility. The company aims to differentiate by giving AI the telecom-specific context, operational access, and controls needed to turn reasoning into trusted OSS outcomes. Lumen adopted AI Studio and pre-built OSS agents in February 2026 for device modeling, inventory reconciliation, and digital-twin insight, and at MWC 2026 Blue Planet demonstrated "Agents-as-a-Service," a supervisor and six specialist OSS agents on Amazon Bedrock Agents and AgentCore, a cloud-delivered path alongside its on-premises "bring your own AI" model.

The most significant 2026 addition is Configuration and Change Management (CCM), launched in June, which extends this approach by bringing network configuration itself under centralized, automated governance. CCM provides configuration visibility and trusted baselines, drift monitoring, compliance validation, controlled change workflows, auditability, and integration with AI Studio and its agentic framework. Together with Blue Planet's data fabric, which provides trusted network and service context, CCM strengthens Blue Planet's ability to connect context, decisioning, execution, and feedback in a governed operational loop. The company showcased a service lifecycle Level 4 architecture and two TM Forum Catalysts at DTW Ignite with operators including KDDI, LG U+, Orange, and Vodafone. Blue Planet was named a Leader in the IDC MarketScape: Worldwide Telecom Autonomous Network Operations 2026.

©2026 AvidThink LLC. All Rights Reserved.

Ericsson

Ericsson supplies mobile network infrastructure and software to service providers and, increasingly, enterprises; its carrier automation platform, Ericsson Intelligent Automation Platform (EIAP), covers telco RAN to core. Meanwhile, Ericsson Enterprise Wireless Solutions (formerly Cradlepoint) offers its cloud-based management solution, NetCloud. Our 2025 profile led with ANA, NetCloud's agentic assistant, but the major update this year is on the carrier side, where EIAP has matured. By June 2026, EIAP had more than 100 members and 100 rApps (such as power-saving apps). Customer AT&T runs third-party rApps in production and has reciprocal membership in Nokia's RAN System and Management Orchestration marketplace.

The agentic additions are currently at the qualification stage: Agentic rApp-as-a-Service on AWS Marketplace (Vivo Brazil testing), the AT&T/Aira intent-to-rApp pipeline, and extending EIAP to the core. TM Forum validations help bolster Ericsson's credentials in this market: NBN Co's Level 4.0 for one fixed-wireless scenario and True Corp's Level 4 service assurance both run on Ericsson intent-based operations.

Huawei

Huawei's carrier business markets its Autonomous Driving Network (ADN) framework with domain solutions for wireless, core, optical, and IP, and Xinghe Intelligent Network for enterprises. Customer references include China Mobile, LG U+, and MTN. Our 2025 profile covered the "Agentic Core" vision; the major update is the move from single-scenario Level 4 to industrializing cross-domain agent collaboration. At MWC, Huawei launched "AN L4 Phase 2" (network-element-level awareness, domain-agent collaboration, and the A2A-T protocol), saying single-scenario ADN was commercially deployed on 130+ networks by end-2025. Huawei's open-source contribution, A2A-T, formalized as a TM Forum standard, moved into the Linux Foundation Networking OpenAN project in June. It was co-launched with China Mobile, Huawei's flagship customer for advanced autonomy: TM Forum's March 2026 report records China Mobile's NOCs moving from a self-assessed 3.2 to Level 4 with more than 30% less back-office staffing and about 30% lower MTTR.

Netcracker

Netcracker, NEC's OSS/BSS and services subsidiary, closed its acquisition of CSG Systems in May 2026 (significantly growing its client base). Netcracker customers include T-Mobile, KDDI, Swisscom, Rakuten Mobile, and du. A year ago, we noted it as an OSS vendor adding agentic features; the major update is scale before autonomy. Its Agentic AI Solution ships 60+ pre-built agents, an AI Trust & Control layer, A2A support, and MCP servers embedded in Resource Inventory, Service Orchestration, CPQ, and Revenue Management. Swisscom won the FutureNet World 2026 Autonomous Network Operations award for a Domain Automation blueprint built with Netcracker. In that customer example, deployment time was cut by more than 60% and development cost by 40%. Netcracker appears in IDC's 2026 Telecom Autonomous Network Operations MarketScape.

©2026 AvidThink LLC. All Rights Reserved.

Nokia

Nokia positions itself as an enabler of autonomous networks for the AI era across fixed, mobile, IP/optical, and data-center domains. Our 2025 profile described its Sense-Think-Act vision; this year's major update is that the vision has become domain-specific agentic frameworks built into each platform. Nokia is aiming for "glass box autonomy" as its organizing principle, focusing on transparent, explainable agent behavior.

In June 2026, Nokia introduced a trust-based agentic AI framework for its Network Services Platform (NSP), deployed by more than 1,000 operators: agents reason over a live network view of the WAN and take guided actions within operator-defined policies, starting with a Troubleshooting Agent, with MCP access for external agents and availability by end-2026. The stack is data curated through an ontology, an agent-and-context framework (lifecycle, memory, RBAC, and validation), model choice, and an Agentic Command Center for natural-language interaction. In the data center, its Event-Driven Automation (EDA) AIOps provides human-gated agentic remediation supported by digital-twin dry-run and rollback capability. Nokia's IT department has migrated its six largest data centers to run on EDA, resulting in significant uptime improvements.

Nokia's agentic work spans every domain it serves: bring-your-own-LLM agents for home and broadband networks, a pre-built Agent Library, and a multi-cloud model that gives operators a choice of runtime. Nokia has partnered with AWS on Amazon Bedrock (with Orange and du as customer examples) and with Google on six Gemini-powered agents in Nokia Assurance Center.

Rakuten Symphony

Rakuten Symphony sells the cloud, RAN, and OSS stack that Rakuten Mobile runs to operators such as 1&1 and AT&T, and its anchor customer provides the most candid autonomy data in the market. In February 2026, Rakuten Mobile received TM Forum Level 4 autonomy validation for RAN energy-efficiency optimization on a live Open RAN network, achieving about 20% energy savings. Rakuten's platform uses LLMs only for hard-to-resolve situations and manages token costs tightly. In Rakuten Mobile's stack implementation, guardrails sit at the point of execution as policy-as-code.

©2026 AvidThink LLC. All Rights Reserved.

Ribbon Communications

Ribbon Communications delivers AIOps and automation platforms with agentic AI technologies, market-leading voice communications software, IP routing, and optical networking to create better experiences, reduce operational costs, and achieve scalable growth. Our 2025 profile highlighted the company's AIOps and automation platform, Acumen, at launch. This year's major update is that Acumen is now deployed with multiple customers and features a more defined agentic architecture. Acumen stands out for its ability to process real-time network operational data at scale across Layers 0–7 for cross-domain observability and ML-based insights. Further, Acumen Builder gives customers an intent-based environment for creating agents and automated operational workflows. At the OFC 2026 optical networking conference, Ribbon also introduced the Muse AI agent for its IP routing and optical networking elements, which answers natural-language questions about the network, with Acumen-on-Muse paving the way for "self-driving" IP optical operations.

In February 2026, Ribbon signed an AWS Strategic Collaboration Agreement, bringing its cloud-native voice portfolio to Amazon EKS and AWS Marketplace, with a Fortune 500 enterprise already in production. The relationship is expanding to include the Acumen platform, which uses Amazon Bedrock for intelligent network operations and Amazon Bedrock AgentCore for agentic workflows, including the Acumen NOC Agent to help with troubleshooting and planning remediation steps. Human approval remains the default, but autonomy can be granted selectively.

Totogi

Totogi, the TelcoDR-backed challenger to the OSS/BSS incumbents, leads with its telco ontology as a core differentiator. It describes its ontology as a guardrail and an executable knowledge graph over existing BSS, OSS, and network systems. Customers include StarHub, 2degrees, and Zain Sudan. The company claims that "nearly 10" tier-1 operators use it; a September roundup publicly lists seven deployments and cites a 97% cut in alarm noise at an anonymized sub-Saharan African tier-1 operator.

APPENDIX | E. Horizontal AI and Workflow Platforms

IBM

IBM positions IBM Network Intelligence (NI) as "network AI that works, not just watches." IBM NI can be operated as SaaS, on-premises, or hybrid. It pairs IBM's Granite time-series foundation models with reasoning agents for telco and enterprise networks, and is sold alongside watsonx and IBM Concert. Our 2025 profile described the launch architecture; this year's major update is a much more specific account of the governance harness around the model. IBM describes a "harness over model" design. Its machine-learning TSPulse models pre-process metrics before the reasoning model gets involved. IBM NI also leverages MCP (with RBAC controls) for real-time context. IBM Concert Workflows are invoked as deterministic APIs (as in the architecture described above). watsonx Orchestrate generates evaluations using watsonx as judge (AI as judge), and instead of a sprawling ontology, IBM favors graph-guided investigation. From a capabilities standpoint, IBM sequences Day 2 assurance first, Day 1 planning second, and will address Day 0 tasks in the future.

©2026 AvidThink LLC. All Rights Reserved.

NVIDIA

NVIDIA's telecom message is "building the AI-native telco" across AI-RAN, sovereign AI infrastructure, and "AI-powered operations"; its relevance to this report is the runtime it is assembling for telco network agents rather than GPUs or fabrics. At DTW Ignite 2026, it presented NeMo Safe Synthesizer and Anonymizer for synthetic telecom data, NemoClaw blueprints for long-running agents with policy guardrails, OpenShell as an open-source sandboxed runtime enforcing policy at the moment of action, and NV-Tesseract time-series anomaly models, with the principle that humans keep control of policy and agent behavior must be auditable. Its partner, SoftBank, fine-tunes its Large Telecom Model on NVIDIA platforms using synthetic data; ServiceNow runs Project Arc NOC agents inside OpenShell; and Gluware's Titan MCP server targets OpenShell as a host. It remains unclear whether network operators will adopt OpenShell as a common agent sandbox.

ServiceNow

ServiceNow sells Telecommunications Service Management as the platform connecting telecom operations "from the network to the customer," with customers such as Bell, BT, and Verizon Business, and at MWC 2026 framed its telecom agents as "Autonomous CRM" governed by its AI Control Tower. This matters to network operations because the ticketing and workflow layer, where most agentic NetOps products terminate, is where it began supplying its own agents in 2026: a network-incident resolution agent coordinating sub-agents, a correlation monitor, next-action and remediation agents, and service test-and-repair agents. These agents create cases, work orders, and repair tasks rather than touching network elements. For enterprises, ServiceNow ITSM and its CMDB are often already the ticketing and change-approval system of record, so enterprise NetOps agents from other vendors frequently open tickets and request change approval there. ServiceNow has greater ambitions with Project Arc, an autonomous agent in early preview that runs inside NVIDIA's OpenShell sandbox with zero default permissions; at DTW Ignite, the two companies described extending it into telecom NOCs with humans involved only when judgment is needed. No telco is yet named as using the telecom agents or Arc.

—

©2026 AvidThink LLC. All Rights Reserved.



AvidThink, LLC
1900 Camden Ave
San Jose, California 95124 USA
avidthink.com

©2026 AvidThink LLC. All Rights Reserved.

This material may not be copied, reproduced, or modified in whole or in part for any purpose except with express written permission from an authorized representative of AvidThink LLC. No part of this work may be used or reproduced in any manner for the purpose of training artificial intelligence technologies or systems. In addition to such written permission to copy, reproduce, or modify this document in whole or part, an acknowledgment of the authors of the document and all applicable portions of the copyright notice must be clearly referenced.

Revision A