



Agentic AI transforms operations in IP networks

How the NSP Troubleshooting Agent helps network operators transition to guided and faster resolution

White paper

Many network operators rely on reactive troubleshooting for their IP networks. They respond after problems have impacted services and then work to resolve them as quickly as possible. For many, the big challenge is to interpret a flood of data to determine what's causing a problem and how to mitigate it.

The Nokia Network Services Platform (NSP) Troubleshooting Agent addresses this challenge by using agentic AI to shift operators' focus from data analysis to guided resolution. It accelerates troubleshooting by correlating symptoms with topology, configuration changes, historical behavior and services to ensure that operators always understand the problem and have the best remediation options.

Read this paper to learn how the NSP Troubleshooting Agent helps operators resolve network problems more effectively and build a foundation for trusted, human-governed autonomous operations.

Contents

Introduction	3
From reactive troubleshooting to guided resolution	3
Why troubleshooting requires an agentic approach	4
The NSP Agent Framework: A foundation for multiple capabilities	4
Context is everything: The role of ontologies and knowledge graphs	5
Human-governed automation and operational trust	5
Alignment with the TM Forum Autonomous Networks Framework	6
Trustworthy AI through proven operational foundations	6
Why evolving existing controllers matters	7
A platform for continued innovation	7
Conclusion	7
Abbreviations	8

Introduction

The Nokia Network Services Platform (NSP) Troubleshooting Agent is the first artificial intelligence (AI)-based capability delivered as part of NSP’s agentic AI roadmap. It is designed to support one of the most critical and time-sensitive processes in IP network operations: identifying the root cause of an issue, determining its impact on services and proposing the most effective corrective actions.

Unlike analytics tools that rely solely on predefined alarm correlation rules, the Troubleshooting Agent works with NSP’s complete operational context. It considers inventory, topology, alarms, service definitions, configuration history and telemetry data to reason about issues.

When an incident occurs, network operators typically face a flood of alarms, counters and logs. Determining which signals matter and how they relate to one another requires experience and time. The Troubleshooting Agent accelerates this process by systematically correlating symptoms with topology, recent configuration changes, historical behavior and service relationships. The result is a guided troubleshooting experience that helps operators focus on the most likely causes and the most relevant remediation options.

Additionally, NSP provides two ways for operators to avoid issues before they arise. First, it keeps track of alarm trends and can identify network resources that require special attention. Second, it can identify “silent failures” on standby resources that do not immediately affect end users but would become visible if a protection switch occurred.

Importantly, the Troubleshooting Agent is intended to augment existing operational workflows rather than replace them. Operators remain in control of decisions and actions, while AI provides structured reasoning, context-aware insights and consistent guidance that reduce cognitive load and shorten investigation cycles.

From reactive troubleshooting to guided resolution

In many operational environments, troubleshooting remains largely reactive. Operators respond to alarms after service degradation has already occurred, often working under time pressure to restore service as quickly as possible. Even in well-instrumented networks, the challenge is rarely a lack of data—it is the effort required to interpret that data correctly and efficiently.

The NSP Troubleshooting Agent addresses this challenge by shifting the focus from raw data analysis to guided resolution. Rather than presenting operators with isolated alerts or dashboards, the agent frames incidents as structured problem statements: what is affected, what has changed, what is likely related and what actions are available.

For example, when alarms are triggered following a configuration change, the agent explicitly links those events. It can highlight configuration drift, identify dependencies between network elements, and surface services or customers that may be impacted. This reduces the need for manual cross-checking across multiple tools and interfaces, which allows operators to move more quickly from detection to diagnosis.

By embedding this reasoning directly into NSP, the Troubleshooting Agent ensures that recommendations are grounded in live network state and operational constraints. This is particularly important in IP networks, where changes made to resolve one issue can easily introduce unintended side effects elsewhere if context is incomplete or misunderstood.

Why troubleshooting requires an agentic approach

Troubleshooting in IP networks is an iterative process that involves forming hypotheses, validating assumptions, examining evidence and refining conclusions as new information becomes available. This makes it fundamentally different from single, linear tasks such as batch analytics or one-off AI queries.

NSP adopts an agentic AI approach to support this type of reasoning. Instead of relying on a single monolithic AI model, NSP introduces intelligent agents that can pursue goals, maintain context over time, and interact with data and tools in a structured manner.

In the context of troubleshooting, agents can:

- Observe network conditions and changes continuously.
- Correlate information from multiple sources, including topology, configuration, telemetry and services.
- Track the state of an investigation as it evolves.
- Explain how conclusions are reached and what evidence supports them.

This agent based model aligns naturally with the way operators work today. Investigations rarely follow a fixed script. They adapt as new clues emerge. By mirroring this behavior, agentic AI allows NSP to support more natural, flexible and effective troubleshooting workflows.

The NSP Agent Framework: A foundation for multiple capabilities

The Troubleshooting Agent is built on a generic Agent Framework that establishes NSP as a platform for multiple future agent-based capabilities. This framework provides the common services required to create, deploy and govern intelligent agents within the platform.

At its core, the Agent Framework manages:

- Agent lifecycle and orchestration.
- Access to operational data and models.
- Interaction patterns between agents and users.
- Governance, auditability, and safety controls.

By separating the agent infrastructure from individual use cases, NSP ensures that new AI capabilities can be introduced incrementally without the need to redesign the platform. While troubleshooting is the first concrete application, the same framework can support additional agents focused on areas such as control plane assurance, change management, optimization or closed-loop automation.

This approach also supports consistency. Operators interact with different agents through a unified experience, and agents follow the same operational rules and safeguards. Over time, this consistency enables a coherent ecosystem of AI-assisted workflows rather than a collection of disconnected features.

Context is everything: The role of ontologies and knowledge graphs

Effective troubleshooting depends on context. An alarm on a router interface means very little on its own. Its significance depends on where the interface sits in the topology, which services traverse the interface, what changed recently and which customers might be affected.

To capture and reason over this context, NSP uses ontology-based data models and knowledge graphs as a foundational layer for agentic AI. These models explicitly define the relationships between physical and logical network elements, services, configurations, historical events and operational policies.

By representing network knowledge in this structured way, NSP enables AI agents to reason over relationships rather than isolated data points. This has two critical benefits. First, it improves accuracy by grounding AI conclusions in validated operational models. Second, it supports explainability. Agents can trace the way a conclusion was reached and which relationships or events contributed to it.

This approach also helps mitigate one of the key risks associated with AI in operations: erroneous or misleading conclusions. By constraining reasoning to known entities and relationships, ontology-driven AI reduces ambiguity and limits the scope for hallucination or unsupported inference.

Human-governed automation and operational trust

While AI creates new opportunities to increase speed and consistency in network operations, the adoption of AI-driven decision-making in live networks must be approached carefully. Operators remain accountable for service quality, availability and compliance, and any automation introduced into the network must earn trust incrementally.

For this reason, NSP's agentic AI capabilities are explicitly designed around the principle of human-governed automation. The Troubleshooting Agent does not execute corrective actions autonomously by default. Instead, it supports operators by explaining its assessment, highlighting evidence and proposing remediation options that can be reviewed and approved before execution.

This interaction model serves three key purposes. First, it allows operators to validate AI reasoning against their own experience and situational awareness. Second, it provides transparency: Operators can understand why a particular recommendation was generated and which data points contributed to it. Third, it ensures accountability remains clear, even as automation levels increase.

Over time, operators may choose to expand the scope of automated execution for well-understood scenarios, such as recurring configuration issues or known failure patterns. NSP supports this evolution by providing explicit guardrails, policy controls and auditability. This approach allows automation policies to be applied selectively and refined as confidence grows.

Alignment with the TM Forum Autonomous Networks Framework

The TM Forum Autonomous Networks framework provides a widely accepted industry reference model for the progressive automation of network operations. It defines autonomy levels that range from manual operations with decision support to fully autonomous systems capable of self-management.

Within this framework, Level 4 autonomy represents a pragmatic and highly valuable milestone. At this level, systems can detect issues, diagnose root causes and recommend corrective actions with minimal human intervention. Execution remains governed and supervised.

With the introduction of the Troubleshooting Agent, NSP enables operators to move decisively toward Level 4 autonomy for troubleshooting workflows. Agentic AI performs continuous analysis, correlates signals across domains and presents structured recommendations. Operators retain approval authority over execution.

This balance is particularly important in IP networks, where operational environments are heterogeneous and risk tolerance varies between operators and use cases. By aligning with the TM Forum model, NSP provides a clear and standards-based path for operators to advance autonomy incrementally, without forcing disruptive operational changes.

Trustworthy AI through proven operational foundations

Trust in AI systems cannot be achieved through technology alone. It depends on how those systems are designed, governed and integrated into real operational environments. NSP's approach to agentic AI draws on the more than two decades of experience that Nokia has gained in managing IP networks at scale across a diverse global customer base.

This experience is reflected in the way NSP models network resources, enforces constraints, and represents dependencies between configuration, topology, services and customers. NSP's AI agents do not operate on raw or unvalidated data. They reason over operational models that have been refined through years of deployment and real-world feedback.

In addition, NSP benefits from broader Nokia expertise in trustworthy AI design. Collaboration across Nokia, including contributions from Nokia Bell Labs, informs NSP's approach to security, explainability and resilience. These principles are embedded into the Agent Framework and the governance mechanisms that surround it.

As AI capabilities evolve, this shared foundation helps ensure that new functionality remains aligned with operator expectations, regulatory requirements and the practical realities of operating live networks.

Why evolving existing controllers matters

The introduction of agentic AI in NSP reflects a broader strategic choice: to evolve existing operational platforms rather than introduce AI as a detached overlay. IP network troubleshooting depends on deep knowledge of how networks are built, configured and operated. This knowledge already resides within the resource controller.

By embedding AI directly into NSP, operators can leverage their existing investments in workflows, policies and operational models. AI becomes an extension of the tools they already trust, rather than a parallel system that must be reconciled manually.

This approach also enables tighter integration between insight and action. Recommendations generated by the Troubleshooting Agent are immediately actionable because they are produced within the same platform that enforces configuration, change control and service definitions. As a result, operators can move more quickly from diagnosis to resolution without context-switching across multiple tools.

A platform for continued innovation

While the Troubleshooting Agent is the first visible application of NSP's agentic AI strategy, it is intentionally designed as a starting point rather than an endpoint. The underlying Agent Framework, ontology based data layer and governance model support the introduction of additional AI-assisted workflows over time.

Potential future applications include proactive assurance, intelligent change validation, optimization guidance and closed loop automation for selected scenarios. Each of these capabilities can be introduced incrementally, building on the same foundational architecture and interaction patterns established by the Troubleshooting Agent.

By treating agentic AI as a platform capability rather than a one off feature, NSP positions itself to evolve alongside operator needs and industry maturity, supporting a sustained transition toward more autonomous and resilient network operations.

Conclusion

AI is moving from experimentation to operational reality in IP networks. With the introduction of the NSP Troubleshooting Agent, NSP delivers a practical, trust-focused application of agentic AI that directly addresses the need for faster and smarter problem resolution, one of the most critical challenges operators face today. It provides several key capabilities that help operators make their networks more efficient and resilient:

- The Troubleshooting Agent improves troubleshooting efficiency by correlating alarms, topology, configuration changes, services and customer context.
- Agentic AI enables iterative, context-aware reasoning that aligns with real troubleshooting workflows.
- Ontology-based data models and knowledge graphs ground AI reasoning in structured, validated operational knowledge.
- Human-governed automation supports trust, accountability and incremental adoption.
- Alignment with the TM Forum Autonomous Networks framework provides a clear path toward Level 4 autonomy.
- NSP's Agent Framework establishes a foundation for continued AI innovation beyond troubleshooting.



Together, these capabilities reinforce NSP's role as a central platform for modern IP network operations, one that combines deep operational context, structured intelligence and practical automation to help operators work more efficiently and confidently in increasingly complex environments.

Abbreviations

AI	artificial intelligence
IP	Internet Protocol
NSP	Network Services Platform
OAM	operations, administration and maintenance

About Nokia

Nokia is a global leader in connectivity for the AI era. With expertise across fixed, mobile, and transport networks, we're advancing connectivity to secure a brighter world.

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

© 2026 Nokia

Nokia Oyj
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: (May) CID215425