HyperFRAME RESEARCH

RESEARCH BRIEF

# Cisco AgenticOps:

## Reshaping Global Resilience through Autonomous Networking and Cross-Domain Intelligence

MARCH 2026

Author:
**Ron Westefall**
VP and Practice Leader for
Infrastructure and Networking

# Executive Summary

Organisational network operations (NetOps) are evolving from just managing devices to supervising a web of distributed, AI-driven capabilities. Existing operational frameworks are failing because they were built for static environments and linear growth. Today's NetOps complexity is exponential. Operations now span millions of ephemeral cloud entities and billions of OT/IT edge devices that generate telemetry volumes that overwhelm human cognition. A war room of engineers manually digging through logs is too slow to troubleshoot and fix problems or prevent service collapse. Machine-speed logic is now required to manage network stability, traffic routing and threat mitigation. To survive this shift, organisations must move from passive monitoring to a system that possesses the context to self-correct before human intervention is even possible. That system is Cisco's AgenticOps.

Cisco's AgenticOps is a framework for managing and scaling fleets of autonomous AI agents with varying levels of autonomy across an enterprise. It provides the operational discipline of deploying, managing, monitoring and governing autonomous AI agents within IT environments. The framework enables agents to observe infrastructure, reason about problems and take corrective actions with minimal human intervention. AgenticOps focuses on the full lifecycle of an agent, providing an infrastructure that maintains reliability, safety and accountability. It also provides governance, transparency and controls that allow operational supervision over agent autonomy and execution.

Cisco's AgenticOps is the next evolution of IT Operations (AIOps). It blends agent-first, purpose-built autonomous action with oversight, into a unified multiplayer, human and agents experience. AgenticOps goes beyond AIOps in significant ways:

- AIOps applies Gen AI and machine learning to detect anomalies and correlate related events. Human operators are needed to interpret recommendations, write automation scripts and manually approve or execute fixes.
- AgenticOps goes further, enabling enterprises to safely and reliably deploy AI agents that can independently reason through problems and complete, at machine speed, complex, multi-step tasks across different software and hardware platforms. AgenticOps brings autonomy into IT operations, amplifying human capacity and skill, turning the network from a reactive system to a proactive, self-optimising one by combining reasoning, simulation and closed-loop execution.

In daily operations, this means moving away from a world of endless dashboards and manual ticket routing. With AgenticOps, agents continuously sense and reason across every dependency before an engineer can receive a notification about a network bottleneck and identify the issue. From local devices to the cloud, agents instantly detect and execute fixes. AgenticOps break down silos across operational domains, maximizing the overall service experience.

Critically, an AgenticOps framework must anchor itself in the organisation's existing workflow reality; autonomous operations are not, and should not be, instantaneous. Agents within the framework provide specific remediation blueprints for human approval and interfacing in a unified multiplayer experience. This enables teams to build trust in the system and delegate authority incrementally, shifting the IT lifecycle from a state of constant manual intervention to one of high-trust, automated oversight.

This Research Brief makes three core points:

1. First, AgenticOps is necessary because the complexity of modern digital infrastructure has reached a tipping point where autonomous, machine-speed execution is essential to closing the gap between real-time telemetry and effective remediation.
2. Next, Cisco's AgenticOps framework transforms the network into a silo-free, autonomous environment by deploying an agentic workforce of specialised agentic capabilities focused on troubleshooting, optimisation, and validation, that use cross-domain telemetry and 40 years of codified expertise to act as collaborative partners alongside IT teams.
3. Finally, we assess why Cisco is the superior provider of AgenticOps for NetOps because of its unique, deeply integrated, and multi-layered AI architecture, which delivers streamlined operations and accelerates time to value through complete cross-domain implementation. We conclude with a practical what to do next plan that helps decision makers turn strategy into phased execution.

# The Strategic Imperative: Why AgenticOps is Necessary

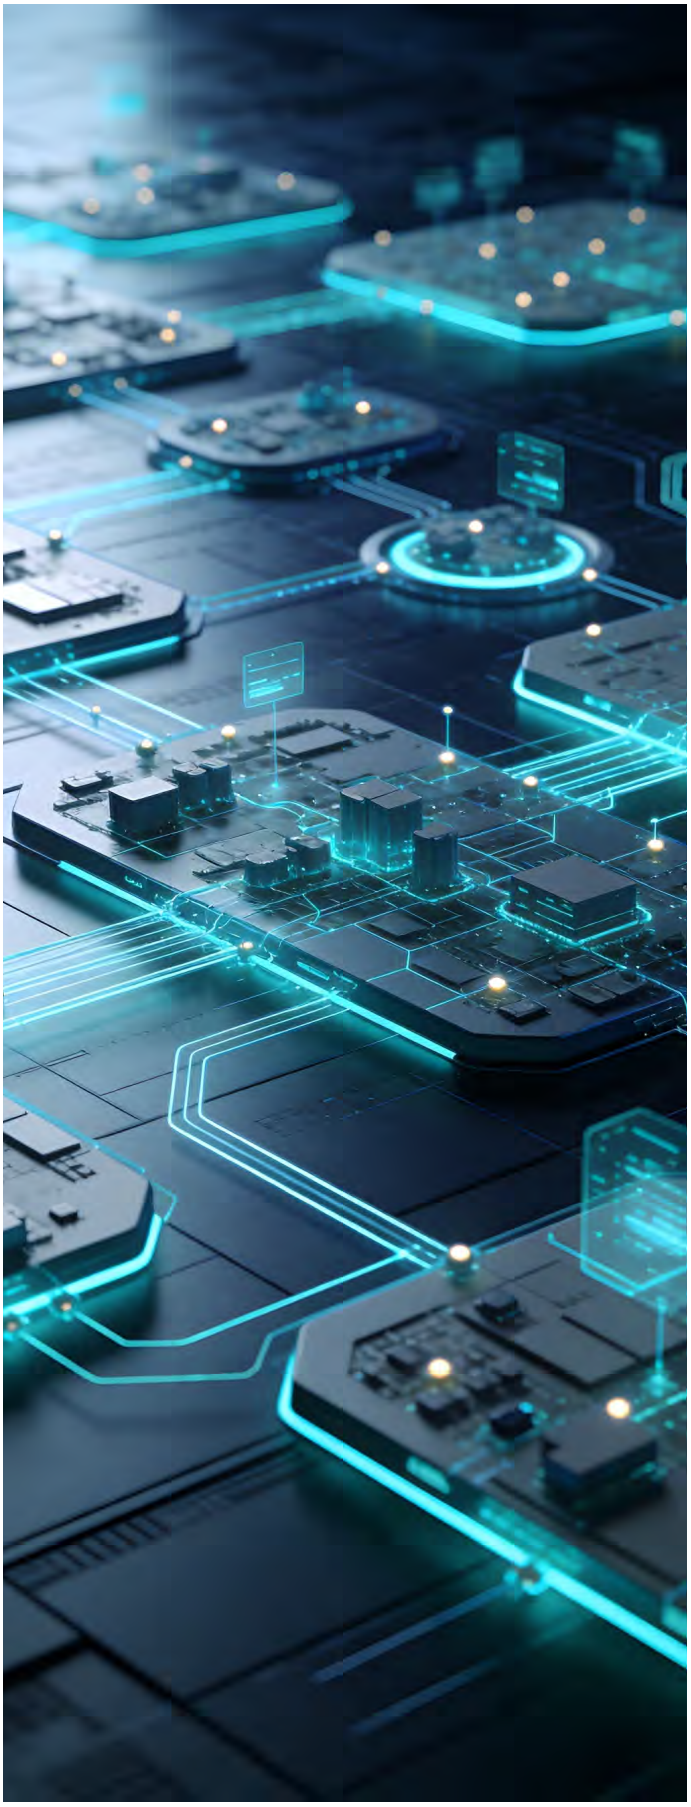

## AgenticOps: Traditional AIOps Is Not Enough

Traditional AIOps solutions have a functional ceiling. First, there was machine learning that could spot anomalies and generate alerts. Next came Gen AI that could explain the issue and recommend fixes. Then came AIOps platforms that are designed to collect telemetry, correlate events, and improve alerting within individual functional silos like networking or security. While AIOps helps to reduce noise, it doesn't change the operating model; humans remain the central engine for issue resolution, manually bridging the gaps between disconnected domains and insights.

The reality of modern infrastructure is that outages and performance degradations rarely originate in a single layer. Instead, they emerge from complex interactions between the network, cloud, security protocols, and application stacks. Because traditional AIOps remain largely siloed and observational, cross-domain issues force teams into manual correlation and war rooms. AIOps detects anomalies in the system but leaves the reasoning to humans. This reliance on human-led troubleshooting leads to alert fatigue and prolonged Mean Time to Resolution (MTTR).

For example, while AIOps is confined to reacting to structured data and threshold-based alerts, such as flagging a traffic spike with a generic High CPU notification, AgenticOps uses LLM-powered, tool-augmented planning to navigate the entire operational stack with human-like reasoning. Instead of merely alerting an engineer, an autonomous agent can investigate the root cause by querying marketing workflow visualisations for active sales, retrieving specific promotional scaling runbooks from internal documentation repositories, executing the necessary scripts to scale resources, and finally summarising the entire resolution in a team messaging channel. This shift moves the system from a passive observer of historical patterns to an active participant capable of cross-platform research and end-to-end task execution. In short, AIOps is a model that cannot scale to meet today's digital NetOps demands.

## AIOps vs. AgenticOps: Key Differences in Operational Impact

| NetOps Activity | Traditional AIOps | AgenticOps |
|---|---|---|
| **Monitoring & Oversight** | Engineers manually monitor dashboards and respond to alerts. | Agents monitor the system constantly; humans move to a supervisory role. |
| **Change Management** | System changes require manual execution and human eyes on glass. | Agents plan, simulate, and execute changes autonomously within safe parameters. |
| **Troubleshooting** | Human-led investigation where engineers dig through logs to find the root cause. | Agent-led diagnosis where agents identify issues, reason solutions and take action with or without humans providing validation. take action with or without humans providing validation |

**Source: HyperFRAME Research**

AgenticOps draws a hard line by prioritising intelligent action. It anchors the concept of a self-healing environment in real-world execution, using agentic AI to predict failures and trigger corrective measures before a performance degradation or worse can propagate across the network. By moving from passive monitoring to machine-speed execution, organisations can prioritise neutralizing infrastructure debt as well as reduce and eliminate the operational downtime that occurs when human-led war rooms simply cannot keep up with the pace of digital disruption and failure.

To achieve enterprise readiness for AgenticOps, organisations must first:

1. Establish a governance framework that defines clear accountability, auditability, and human-in-the-loop guardrails for autonomous agents.
2. Prioritise a modular, data-driven architecture, using a crawl, walk, run approach, that ensures high-quality data fabric and low-latency integration across existing workflows and systems.
3. Adopt an end-to-end AgenticOps framework that combines domain-specific intelligence with unified cross-domain telemetry and human-in-the-loop governance to ensure readiness, address skill shortages and ease adoption.

# From Tool to Teammate: The Rise of Cisco's Agentic Workforce

AgenticOps evolves AI into a sophisticated workforce of autonomous digital peers. By embedding purpose-built CCIE-level logic into the platform, Cisco provides always-on capabilities that move beyond basic automation into active collaboration. Its agentic capabilities leverage deep operational intelligence to proactively manage the network, offering flexible levels of autonomy that range from human-triggered investigations to continuous, deep-thinking environmental assessments. This shift marks a transformative transition to AI as a collaborative partner that runs the network alongside human operators, forming an agentic partnership, where AI and human expertise coalesce to orchestrate the network.

Cisco recently unveiled the next evolution of its AgenticOps framework that prioritises simplified operations and AI-driven solutions that advance the shift toward a network that is autonomous and silo-free. At the forefront of this evolution is the introduction of expanded Troubleshooting, Optimisation, and Validation agentic capabilities designed to run network environments with minimal manual intervention. These agentic capabilities, alongside expanded AI enhancements, are integrated into commonly used IT interfaces including Cisco's AI Assistant, Agentic Workflows, Cisco Cloud Control platform, AI Canvas (coming soon) as well as third party applications.

### Cisco's Agentic Workforce: Agentic Capabilities

| Agentic Capability | Primary Charter | Key Functionality | Advanced Tools |
|---|---|---|---|
| **Autonomous Troubleshooting** | Incident Ownership (Detection to Resolution) | Apply reasoning from telemetry to root cause, validating multiple hypotheses simultaneously and executing deterministic remediations with CCIE-grade precision. | AI Packet Capture (correlating thousands of signals in real-time). |
| **Continuous Optimisation** | Continuous Improvement (Performance & Efficiency) | Continuously maintains user experience by autonomously tuning RF, QoS, path, and control planes with a live understanding of end-to-end network conditions. | AI Configuration Recommendations (proactive tuning for predictability). |
| **Trusted Validation** | Change Safety (Intent to Outcome) | Risk-aware agentic assessments validate network changes against live topology, configuration, and telemetry, including identifying impact and blast radius. | Impact Modelling (reasoning how changes propagate through the system) |

**Source: HyperFRAME Research**

The first of this agentic workforce is the Autonomous Troubleshooting, whose primary charter is incident ownership from detection to resolution. These capabilities investigate, reason, and resolve problems across domains by correlating telemetry from networking, security, and internet layers in real time. Rather than relying on hypotheses, they analyze signals at scale to identify root causes and recommend or execute remediation steps. This capability is bolstered by new tools such as AI Packet Capture, enabling the agent to process thousands of signals simultaneously to provide credible evidence and explanations in minutes.

Supporting the health of the network are the Continuous Optimisation agentic capabilities, digital operators focused on the continuous improvement of performance and efficiency. These agentic capabilities act proactively by detecting configuration drift and predicting degradation across wireless, switching, and WAN environments before users are impacted. Instead of waiting for a support ticket, they recognise emerging risk patterns and tune the environment within defined guardrails. These capabilities are receiving expanded skills, including AI Configuration Recommendations designed to proactively ensure predictable network performance.

The next actors in this agentic workforce consist of Trusted Validation agentic capabilities, which are tasked with making network changes safer and more predictable. These capabilities model the potential impact and blast radius of a change before it occurs, revealing hidden dependencies and downstream risks. Once a change is implemented, the agents automatically verify the outcome and learn from the results to improve future actions. By reasoning how changes propagate through the system, these capabilities ensure that every update moves the network closer to its intended state without unintended consequences.

Cisco's agentic capabilities are designed to be flexible, operating in different modes based on the required level of autonomy:

1. On-Demand capabilities are human-triggered and focused on finite goals, such as resolving a specific client issue.
2. Ambient Agentic capabilities are always-on, pursuing continuous goals like Wi-Fi optimisation or policy drift correction.
3. Deep-Reasoning mode handles long-running, complex objectives such as environment-wide validations or large-scale planning.

Over time, every agentic capability in the Cisco ecosystem is expected to span all three modes, providing a comprehensive and adaptive support system for IT teams.

# From Telemetry to Trust: The Foundational Architecture of Cisco AgenticOps

Cisco's agentic framework is built upon four strategic architectural pillars, beginning with a foundation of cross-domain telemetry. We find that Cisco's agentic technology uses one of the industry's most extensive telemetry sets, spanning campus, branch, WAN, and data centre environments, as well as security layers and application paths. This comprehensive, end-to-end vantage point enables agents to reason and execute actions across the entire system, ensuring that insights are informed by the full scope of the network rather than isolated data points.

The second pillar focuses on Ensemble Models and Codified Cisco Expertise, which transforms general AI into specialised operational intelligence. By combining frontier and foundational models with purpose-built models like the Deep Network Model, Cisco grounds its agentic capabilities in Cisco Certified Internetwork Expert (CCIE) knowledge packs and human-curated runbooks. We see this tapestry of models as enabling the system to choose the right tool for the task, whether the situation requires rapid response, granular troubleshooting, or solving complex architectural challenges.

To ensure these actions are executed safely, the third pillar integrates tools, guardrails, and trust by design. Cisco's agentic capabilities do not act blindly; they use MCP and APIs to gather data while operating within explicit guardrails and customer-defined approval models. To maintain transparency, every action is accompanied by clear reasoning, evidence, and a full audit trail. For Cisco, trust is treated as a fundamental system property rather than just a feature, ensuring that IT teams remain in control of the automated execution.

Finally, the architecture is rounded out by multi-modal interaction, recognising that IT teams require flexibility in how they engage with AI. Rather than forcing a single interface, Cisco surfaces its agentic capabilities through various touchpoints, including existing GA dashboards and the AI Assistant, as well as emerging platforms such as the AI Canvas, messaging systems, and third-party tools such as ServiceNow. By integrating with event-driven alerts and third-party interfaces, the system adapts to the established workflows of modern IT organisations, meeting teams wherever they are most productive.

**Cisco Agentic Framework for NetOps: The Four Strategic Pillars**

| Pillar | Core Objective | Key Capabilities | Why It Matters |
|---|---|---|---|
| **Cross-Domain Telemetry** | Visibility & Context | Spans Campus, Branch, WAN, Data Centre, Security, and Application paths. | Eliminates silos by allowing agents to reason across the entire system. |
| **Ensemble Models & Expertise** | Operational Intelligence | Combines frontier models with a purpose-built Deep Network Model and CCIE knowledge packs. | Moves beyond general AI to specialised, expert-level troubleshooting logic. |
| **Tools, Guardrails, & Trust** | Safe Execution | Uses MCP/APIs within explicit guardrails; provides reasoning and audit trails. | Ensures AI doesn't "act blindly" and keeps human operators in control. |
| **Multi-Modal Interaction** | Workflow Integration | Accessible via Dashboards, AI Assistants, AI Canvas, and messaging. | Adapts to how IT teams work rather than forcing a single interface. |

**Source: HyperFRAME Research**

# The Cisco Edge: Securing the Future of Autonomous Networking through AgenticOps

Cisco's AgenticOps framework stands out by tailoring its intelligence to address distinct customer needs, which includes:

1. **Deep Network Model:** a purpose-built Large Language Model (LLM), trained on over four decades of Cisco intellectual property, interprets the underlying logic of network operations, and delivers real-time insights by directly integrating with live telemetry.
2. **Cross-Domain Scope:** unified visibility spanning all Cisco workloads from Network, Security and Collaboration.
3. **Data Substrate:** a massive data lake underpinned by the harnessing of Splunk (security and log), ThousandEyes (internet and SaaS visibility), and Meraki (cloud networking) capabilities.
4. **Connected Agent Intelligence:** the integration and coordination of network, security, and collaboration technologies across Cisco's platforms and solutions enabling agents and tools to seamlessly share information and work together in real time.

Cisco's Deep Network Model empowers agents to execute with broad operational intelligence. While rivals often operate within silos of wireless diagnostics or data centre automation, Cisco's agentic capabilities draw from a massive data lake underpinned by the harnessing of Splunk, ThousandEyes, and Meraki capabilities. This allows the AgenticOps framework to move beyond simple pattern recognition, using CCIE-level logic to execute complex, autonomous actions that span the entire enterprise stack from the user's device to the cloud application. This model is further refined by over 3,000 reasoning traces, which are expert-derived logic paths that ensure the AI mimics professional human diagnostic steps rather than relying on statistical guesswork.

The Splunk platform data moat differentiates Cisco's AgenticOps framework from the offerings of HPE Juniper and Arista. Cisco's agents operate with a cross-domain context that spans the entire packet lifecycle. Cisco combines real-time telemetry from Meraki, ThousandEyes, and Splunk's extensive security and log data fabric into a unified AI Canvas, which provides agents with the visibility to track a user's experience from a home Wi-Fi connection, across the public internet, and deep into the cloud-native application backend.

This telemetry enables Cisco to move beyond the siloed, single-domain diagnostics typical of its rivals. In a Cisco-powered environment, an autonomous agent can correlate a dip in application performance with a specific security update or a regional ISP outage, allowing it to pinpoint root causes that would remain invisible to vendors focused strictly on the internal network. This comprehensive data substrate ensures that when a Cisco agent acts, such as rerouting traffic or adjusting security policies, it does so with a total understanding of the business impact, offering a level of full-stack operational intelligence that is out of reach for Arista and HPE Juniper.

Finally, Cisco uses cross-domain Connected Agent Intelligence, that enables agents to work together across networking, security, and collaboration tools such as Webex, to resolve issues without manual handoffs. This is a level of integration that competitors such as HPE Juniper or Arista, with narrower portfolio focus, cannot match. This horizontal integration ensures that performance issues are mitigated across different departments and technology stacks before end users experience disruptions. For example, a network agent identifying latency can automatically coordinate with a Webex agent to improve call quality or trigger a security agent to isolate a compromised device, all without manual interventions or handoffs between IT silos.

Cisco prioritises AI safety and trust through robust explainability and built-in safeguards. To ensure reliability, the system provides transparent reasoning for every agent action, enabling:

- **AI hallucination mitigation:** where every autonomous action is accompanied by a Chain of Thought reasoning path that explains the specific data and logic used to reach a conclusion.
- **Rollback automation:** that protects the network from unintended consequences. The framework includes automatic rollback mechanisms that immediately revert configurations if the AI-initiated change causes performance degradation, ensuring high availability even during automated optimisations.

As a result, we find that Cisco holds a unique competitive position as the only vendor capable of bridging the gap between networking, security, and application platforms.

## Competitive Comparison: Cisco vs. Rivals

| Feature | Cisco AgenticOps | HPE Juniper / Arista |
|---|---|---|
| **Intelligence Engine** | **Deep Network Model:** Trained on 40+ years of Cisco IP and CCIE logic. | General-purpose models or siloed AIOps pattern matching. |
| **Scope of Data** | **Cross-Domain:** Unifies Networking, Security, and Collaboration (Webex). | Primarily focused on internal network or wireless silos. |
| **Data Substrate** | **Splunk + Meraki + ThousandEyes:** A massive "data moat" spanning owned and unowned networks. | Limited to proprietary hardware telemetry or specific cloud tools. |
| **Cross-Agent Sync** | **Horizontal Integration:** Agents in one domain can signal agents in other domains to resolve issues without manual handoffs. | Vertical silos; requires manual handoffs between different IT departments. |

**Source: HyperFRAME Research**

# Cisco's Portfolio Advantage: Pioneering the Autonomous Enterprise through AgenticOps

We believe Cisco's main advantage is its transformative AgenticOps framework, which moves network management from manual tasks to an automated, unified digital system. This evolution is led by a specialised team of AI Agentic capabilities for Troubleshooting, Optimisation, and Validation, that act as digital CCIEs to manage everything from incident resolution to change safety. The end-to-end architecture is supported by Cisco's four key pillars for AgenticOps:

1. Extensive cross-domain telemetry
2. An ensemble of expert-informed models
3. Strict trust guardrails
4. Flexible multi-modal interfaces like Cisco's AI Assistant and the AI Canvas

In our view, Cisco is firmly positioned as the reliable partner that organisations need for a successful AgenticOps implementation strategy. Cisco combines its Deep Network Model that ensures diagnostic accuracy, with the massive data fabric of Splunk with telemetry from Meraki and ThousandEyes, to provide a unified view of the network that few competitors can begin to match. Cisco agentic capabilities are highly versatile, operating in On-Demand, Ambient, or Deep-Thinking modes to suit different operational needs and levels of autonomy for each organisation.

With 68% of enterprises planning to replace their foundation models each year[1], Cisco AgenticOps delivers a significant advantage over its competition, providing seamless switching of AI engines without impacting core business logic or accumulating technical debt. Moreover, Cisco's AgenticOps directly addresses the hallucination and security concerns of 90% of our survey respondents[1], by providing traceable, agent-callable APIs and safety guardrails that ensure autonomous actions remain governed and reliable.

In summary, we advise that organisations adopt an incremental path to autonomous network operations. We expect that AgenticOps will be adopted by most enterprises in a crawl, walk, then run approach; beginning with agent-assisted diagnosis and human-approved remediation, expanding into closed-loop execution for well-understood, low-risk actions with clear rollback paths, to more autonomous agent-managed operations as trust in the system grows. Over time, AgenticOps' governance controls, auditability, and blast-radius containment will become just as important as the AI itself.

1 (HyperFRAME Research Lens: 1Q 2026)

# Recommendations for Transitioning to and Evaluating AgenticOps

- **Advance Strategic Transformation with AgenticOps:** Key decision makers, such as CTOs, CIOs, VP of Infrastructure/Operations, VP of Network Operations (NetOps), VP of Security Operations (SecOps), Director of Cloud/IT Strategy, and Chief Architect, need to prioritise evaluating Cisco's AgenticOps framework since it transforms AI from a basic productivity tool into an agentic workforce of specialised agents capable of autonomously troubleshooting, optimising, and validating complex environments at scale. By leveraging forty years of codified networking expertise and unified telemetry across NetOps, SecOps, and cloud domains, the framework significantly reduces operational risk and accelerates time-to-resolution while maintaining human oversight through transparent, trust-based guardrails.
- **Emphasise Comprehensive Autonomous NetOps:** Organisations need to consider adopting the Cisco AgenticOps framework because it offers a comprehensive agentic workforce that can autonomously handle troubleshooting, proactively optimise performance, and validate changes with risk awareness across complex network environments, directly enhancing their teams. By operating in various execution modes, from on-demand assistance to deep-thinking planning, agents use real-time telemetry to resolve incidents and prevent issues before they affect users.
- **Prioritise Trust-Centric Agentic Framework Adoption:** Cisco should be considered as the trusted AgenticOps adviser because its framework draws on forty years of expertise and cross-domain telemetry ensuring autonomous actions are guided by specialised, human-curated operational intelligence. Cisco's AgenticOps framework also prioritises transparency and control by using clear guardrails and multi-modal interfaces, making sure that every agent's action is auditable and aligned with established IT workflows.

## ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programmes.

Our industry analysts specialise in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

## CONTACT HYPERFRAME RESEARCH:

**Steven Dickens**
CEO & Principal Analyst | HyperFRAME Research

**Email Address:**

steven.dickens@hyperframeresearch.com

**Telephone Number:**

+1 845 505 1678

**X: @StevenDickens3**

**LinkedIn: Steven Dickens**

**BlueSky: Steven Dickens**

## CONTRIBUTORS

**Ron Westefall**
VP and Practice Leader for Infrastructure and Networking

## INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

## CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

## LICENSING



## DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.


HyperFRAME
RESEARCH