POV: **Agentic AI**

An agentic AI journey to the Self-Driving Network

How we're delivering on agentic AI's potential with a clear vision and proven results

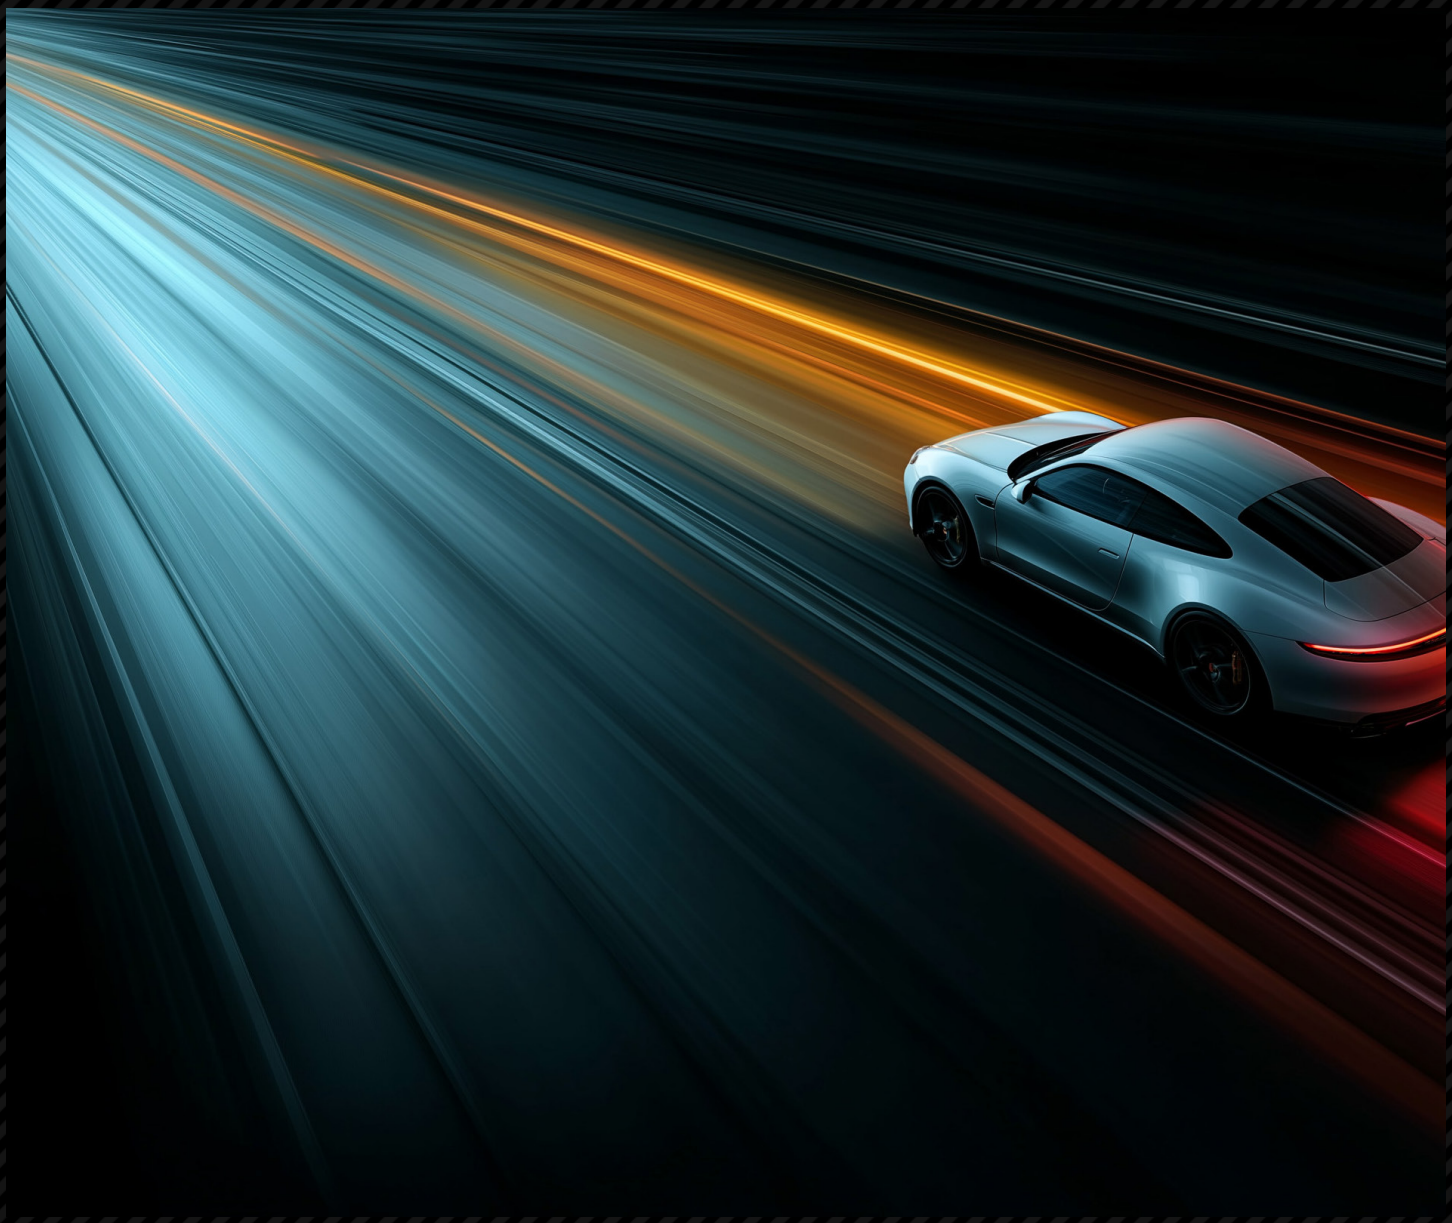

JUNIPER NETWORKS

www.juniper.net

Contents

01 Introduction

Today's IT teams are under constant pressure

They must manage growing complexity with fewer resources, while tight budgets, limited staff, and demanding timelines are the norm.

With Mist™ AI-native networking as a proven foundation, HPE Networking is at the forefront of helping teams meet these challenges—delivering solutions that streamline operations, improve efficiency, and support smarter, more scalable network operations.

Agentic AI has emerged as a powerful tool in the evolution of automation, acting as the catalyst to accelerate the journey to a self-driving network. But agentic AI is not a bolt-on technology that solves all problems. It requires great data along with real-time, AI-native, cloud-native networking. This includes high-quality data, digital twinning, purpose trained large language models (LLMs), models that can predict user experiences, generative AI (GenAI), natural language understanding (NLU), natural language processing (NLP), and an open ecosystem. Together, these elements enable seamless communication across intelligent AI agents, tools, multivendor networking domains, and external platforms.

At HPE Juniper Networking, leveraging agentic AI is the next step in a decade-long journey rooted in our vision of delivering the Self-Driving Network™: an AI-assisted network that helps IT teams optimize network operations and heal itself autonomously. Now, with the emergence of agentic AI, the transition from assisted to fully autonomous operations accelerates—delivering seamless, client-to-cloud experiences with unprecedented intelligence and automation.

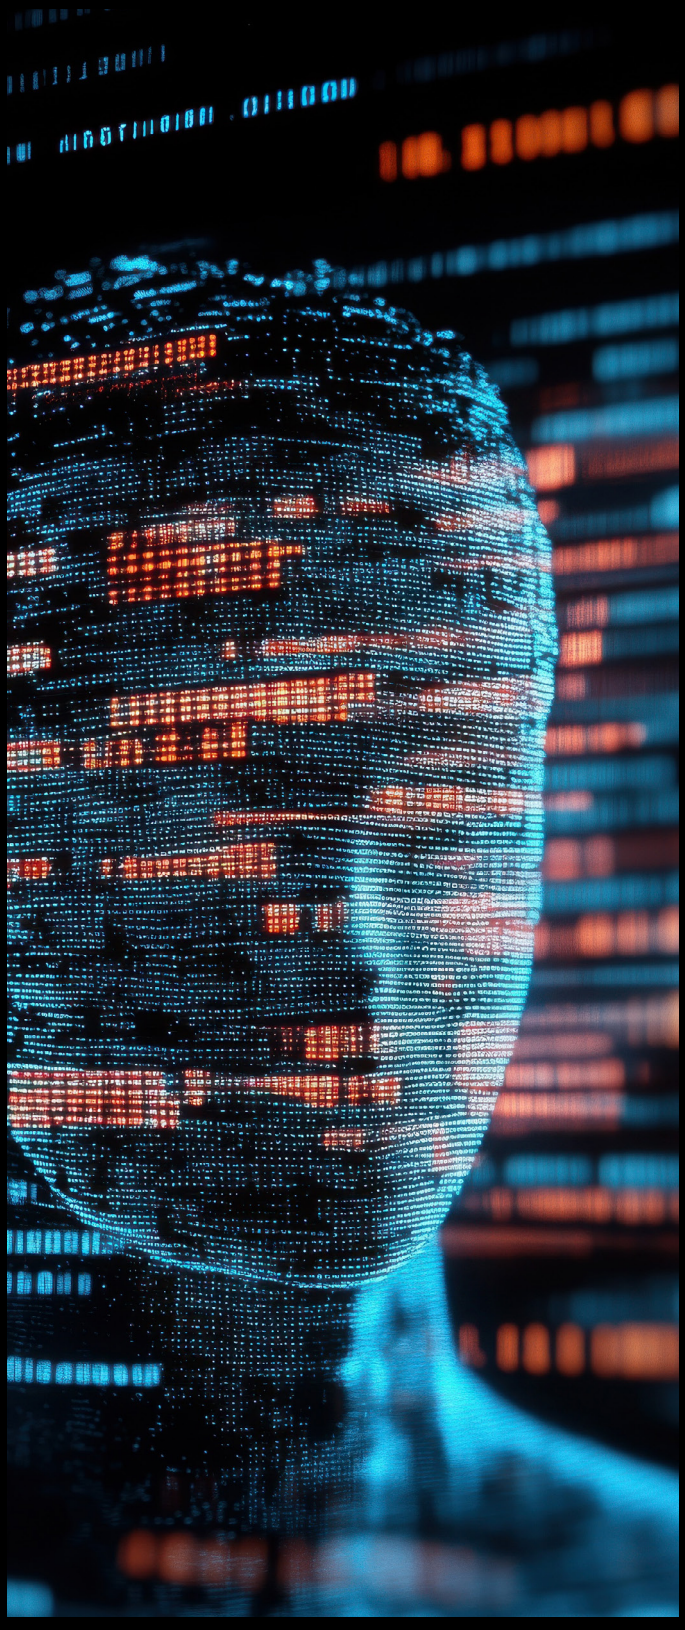

02 Inside agentic AI

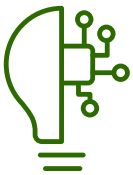

Agentic AI represents the next evolution in automation

AI tools and applications—such as chatbots, LLMs, AIOps, and cognitive search engines—have transformed the way humans interact with technology

By integrating this intelligence across workgroups and resources, agentic AI promises to further accelerate AI innovation by giving agents the agency to make decisions and perform tasks without human intervention. It enables the automation and simplification of complex, labor-intensive workflows such as financial planning, drug development, and cancer research, accelerating efficiency and minimizing human-induced delays.

For service providers, cloud platforms, and enterprises that depend on networks for their core business operations, technologies like digital twins, Large Experience Models (LEMs), and agentic AI will reshape the role of network operations—streamlining processes, reducing manual effort, and driving smarter outcomes.

Agentic AI is a new, nonlinear, nondeterministic programming paradigm. It introduces intelligent agents capable of reasoning, self-learning, and collaborating across both networking and cloud domains. Multiple agents can be loosely coupled, allowing them to operate independently and, when authorized, interact through well-defined interfaces to extend easily to other systems, agents, or tools—such as databases, applications, and web services.

In the networking domain, each agent functions as a domain-specific expert—specialized in areas like switching, routing, or wireless communication. To coordinate their actions, agents are organized within a graph that acts as a blueprint to map how agents interact, plan, and execute tasks. This graph enables orchestrated workflows that empower agents to solve complex problems, adapt dynamically, and autonomously mitigate issues.

Agentic AI decomposes multifaceted challenges into discrete subtasks, delegating them to the most suitable agent or agents. These agents communicate within defined guardrails, using the graph to establish workflows connecting agents and tools. Relevant data is gathered and synthesized by a large language model LLM), enabling intelligent analysis, contextual understanding, and precise resolution—with little to no human intervention.

Agentic AI supports standardized, open protocols to facilitate communication, including the Model Context Protocol (MCP). MCP is optimized for agent-to-API interactions, enabling agents and LLMs to seamlessly access external tools, services, and data sources. While MCP enhances interoperability and scalability, agentic AI does not rely on MCP to function—its outcomes can be achieved independently. As the agentic AI ecosystem matures, its supporting toolkit is expected to evolve, incorporating new protocols and expanding capabilities over time.

We have yet to fully realize the potential value and use cases for agentic AI, but the operational benefits are significant:

- **Speed:** Rapid analysis, synthesis, and resolution of complex problems, minimizing delays and enhancing responsiveness
- **Productivity:** Intelligent, self-learning assistants eliminate mundane tasks, enabling teams to focus on strategic, high-value initiatives and upskill with ease
- **Extensibility:** AI systems evolve and adapt by integrating new agents with specialized capabilities—without disrupting existing workflows—making it easier to tackle complex tasks across industries
- **Innovation:** AI-powered simulations and accelerated prototyping uncover promising models and breakthrough solutions across domains for end-to-end assurance in the networking domain

03 The agentic journey

Accelerating the shift from “assisted” to “self-driving” network operations

HPE Juniper Networking’s bold vision for the Self-Driving Network began with a clean slate—unburdened by legacy architectures built before the AI era

For over a decade, we’ve been steadily advancing toward that vision, evolving the Mist platform with increasingly sophisticated capabilities while continuously enhancing the precision and efficacy of AI-native outcomes.

Agentic AI accelerates the Journey to Self Driving

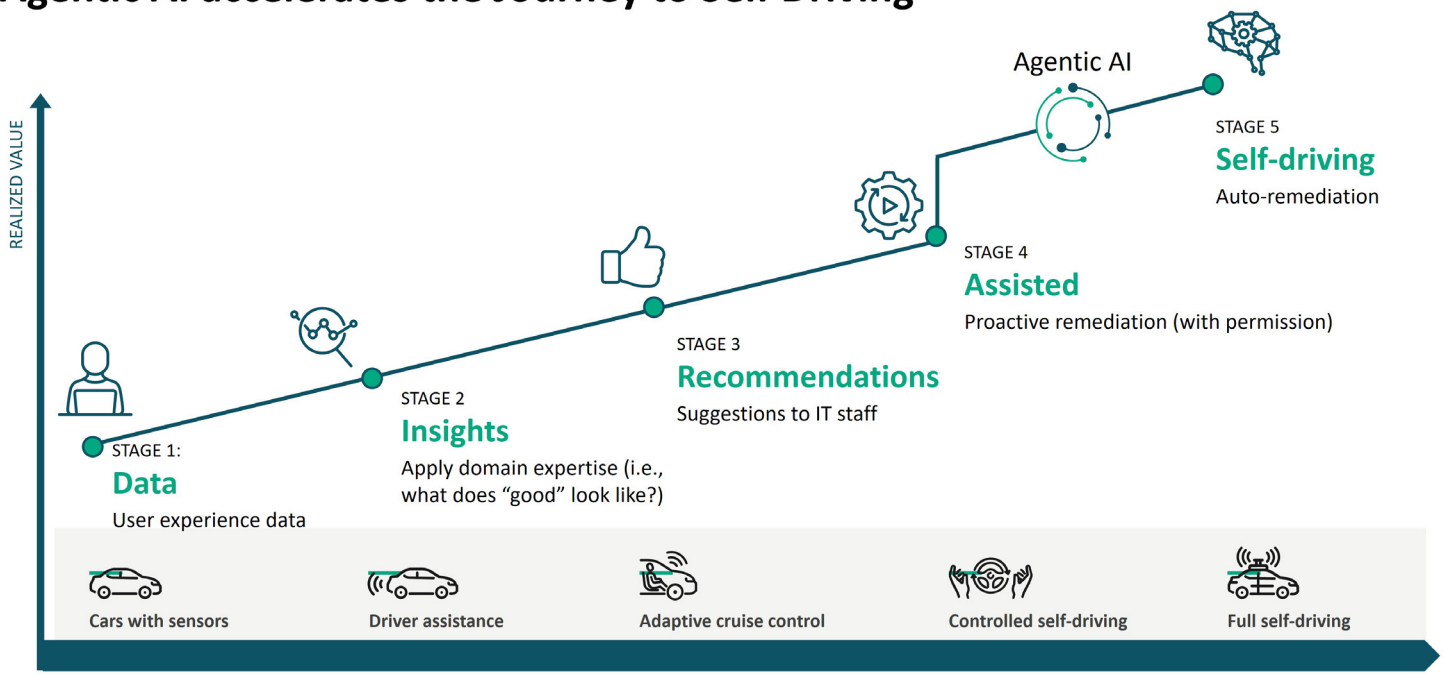

Unlike traditional approaches that focus solely on performance metrics, the Self-Driving Network focuses on predicting and optimizing user experiences. A self-driving network intelligently correlates network analytics—such as loss, latency, and jitter—with user experience metrics at the application level to proactively detect and isolate issues before they impact users. This experience-first model ensures that network operations are not only efficient but also align with user and business needs.

A journey to the Self-Driving Network follows a five-stage evolution built on progressive levels of trust. Much like the evolution of self-driving cars, confidence grows gradually. Drivers first embraced features like automatic braking and lane assist, gaining trust with each step before accepting full autonomy (e.g., hailing a driverless taxi). In the same way, agentic AI earns its place by consistently delivering reliable outcomes—building confidence one action at a time.

Today, many customers are well along this journey, leveraging Mist AI, powered by the Marvis® AI engine. Mist delivers AIOps and self-driving actions for campus, branch, data centers, and WAN, placing HPE Juniper Networking ahead of siloed platforms that are limited to specific domains. Within Mist, AI agents help accelerate the evolution from assisted to self-driving, transforming traditional automation into intelligent, autonomous systems.

Like the adoption of self-driving cars, operations teams need time to build confidence in autonomous capabilities before delegating full control. Mist is already proving value with customers—building trust through evidence-based insights and intelligent recommendations. Guided by human-in-the-loop (HITL) controls, Marvis Actions delivers a range of autonomous functions today, helping teams grow comfortably with increasing levels of automation. As the power of agentic AI continues to expand within Mist, its capabilities will be readily accessible—able to scale in alignment with the trust and confidence it earns.

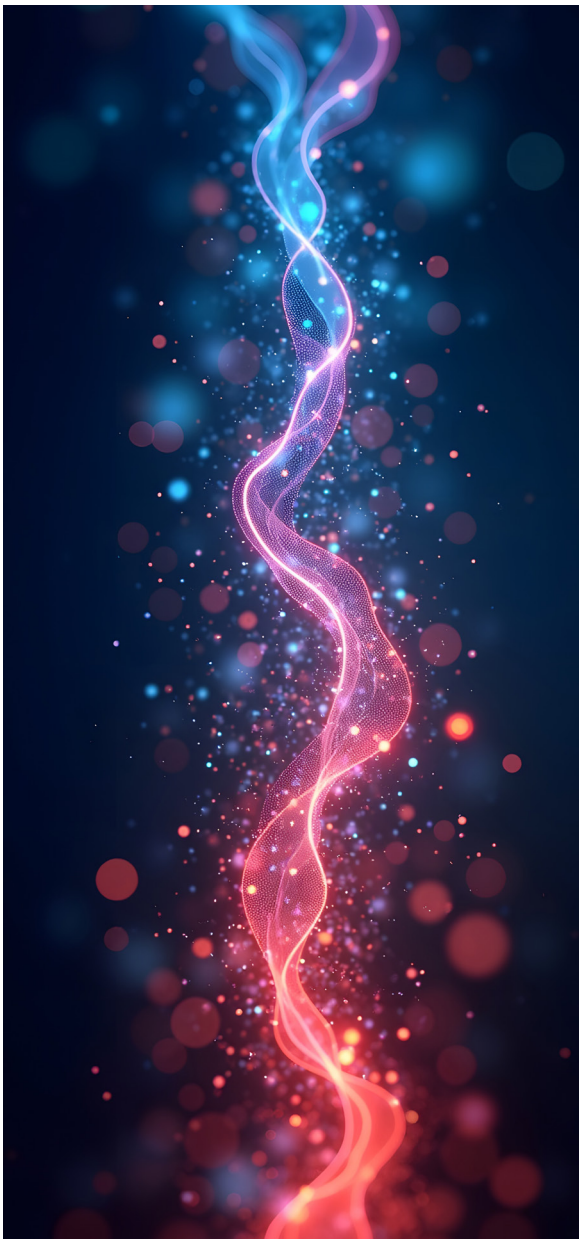

04 A new era in AI

Transforming how IT teams operate and interact

At the core of Mist is the Marvis AI engine and Marvis® AI Assistant

Marvis AI redefines how IT teams interact with and operate their networks. With the integration of AI agents, Marvis AI can reason, collaborate, and act across complex environments, bringing our vision of the Self-Driving Network closer to reality.

Marvis AI Assistant

Marvis AI Assistant includes Marvis Minis, the Marvis Large Experience Model (LEM), the Marvis Conversational Interface, and Marvis Actions. Powered by GenAI and AI agents, each component plays a critical role in interpreting human- or machine-initiated prompts to enable autonomous capabilities

Marvis Conversational Interface

The Marvis Conversational Interface uses advanced LLMs, GenAI, and NLU/NLP to let IT teams ask questions using natural language and receive clear, actionable answers. It understands user intent, engages specialized agents, and orchestrates multistep workflows to diagnose issues, resolve problems, and retrieve documentation. In data center environments, the interface integrates with Apstra's contextual graph database to handle complex queries with a pathway to enable autonomous service provisioning and self-driving data center operations. This combination of conversational intelligence and automation empowers IT teams to operate more efficiently, reduce resolution times, and focus on strategic initiatives that drive innovation.

Marvis Minis

Marvis Minis serve as always-on digital experience twins that continuously simulate user behavior to proactively monitor and test critical network services and applications. By validating performance from client to cloud, they surface issues like misconfigurations or service degradation before users are affected and feed real-time insights into the Marvis AI engine for faster, more reliable operations.

Marvis Large Experience Model

HPE Juniper Networking's Marvis Large Experience Model (LEM) is a specialized model trained on billions of real-world data points to predict user collaboration experience with high precision. Ingesting application data from sources like Teams and Zoom, the Marvis LEM with Shapley looks at the contribution of network features to user experience minutes. Shapley Additive Explanations (SHAP) provides transparent, visual explanations that help IT teams quickly understand and address features and attributes contributing most to root causes.

Marvis Actions

Marvis Actions enables IT teams to shift from reactive troubleshooting to proactive operations by identifying and addressing critical issues—such as DHCP failures, VLAN misconfigurations, and network loops—often before users are affected. It delivers clear, evidence-based recommendations and, for HITL approved scenarios, can autonomously resolve problems like shutting down ports, upgrading devices, or enforcing policy compliance. Every action is validated post-remediation via the Marvis Actions dashboard, ensuring accuracy, maintaining human oversight (HITL) and building trust as teams increasingly rely on Marvis AI's agents to handle routine tasks and focus on higher-value initiatives.

Open ecosystem

Mist is built on a modern, microservices-based cloud architecture with 100% open and programmable APIs, enabling seamless integration with third-party tools, custom applications, and multivendor environments. Open protocols like MCP allow the Marvis LEM and intelligent HPE agents to interface with third-party agents and tools. Acting as universal translators, MCP decouples tools from models, enabling agents to invoke external resources without needing to understand their internal workings. This abstraction streamlines orchestration across domains, accelerating automation and interoperability at scale.

05 Client-to-cloud AIOps

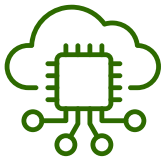

Mist is foundational to HPE Juniper Networking's unified networking strategy

Our goal is to expand AIOps using AI agents across the entire client-to-cloud ecosystem—from wireless access points (APs) to the GreenLake cloud

This approach enables transformative automation that accelerates innovation and delivers business agility at scale. It's a strategy built upon three pillars:

01 AI-native operations

- **The right data:** High-quality, real-time telemetry collected with deep domain expertise to answer experience-first questions
- **The right infrastructure:** A scalable, microservices-based cloud that processes massive amounts of data in real time
- **The right response:** Real-time insights and actions enabling faster root cause identification, smart recommendations, and automated resolutions

02 Comprehensive client-to-cloud portfolio

- **The right networks:** An end-to-end networking portfolio across enterprise, service provider, and cloud provider domains with AI-enabling telemetry and insights
- **The right cloud:** Cloud-native and AI-native hybrid cloud platforms that deliver speed, agility, and scale with cloud simplicity and economics

03 Integral security

- **Unified protection:** Security that protects users, systems, and data across the entire client-to-cloud environment
- **AI-native security:** Protection that continuously learns and adapts to new threat vectors

HPE Juniper Networking stands alone in the industry with a secure AI-native networking platform purpose-built with and for AI. It's paired with a comprehensive client-to-cloud networking portfolio that paves a true agentic path to the Self-Driving Network. Our differentiated mix of networking, security, cloud, and AIOps technologies work in concert to accelerate innovation and deliver measurable business outcomes.

06 Conclusion

The future of networking is self-driving

Agentic AI represents the next major milestone and catalyst in HPE Juniper Networking's journey toward the Self-Driving Network

By enhancing Marvis AI with intelligent agents that reason, plan, and act across complex environments, we're enabling IT teams to operate with greater speed, precision, and proactivity.

Our secure AI-native networking platform unifies operational teams through a shared architecture, enabling faster decisions, coordinated actions, and consistent outcomes—building confidence while streamlining cross-domain collaboration. Trust remains essential in this journey, and we recommend adopting agentic AI gradually, starting with specific use cases and levels of autonomy to balance oversight, value, and outcomes.

With a purpose-built architecture, deep domain expertise, and a unified client-to-cloud portfolio, HPE Juniper Networking is uniquely positioned to lead the industry into the era of agentic networking. As organizations face growing complexity and demand, agentic AI offers a clear, scalable path forward, delivering intelligent automation and measurable outcomes at every stage of the journey to a truly self-driving network.

Why Juniper

Juniper Networks is leading the convergence of AI and networking. Juniper's Mist™ AI-native networking platform is purpose-built to run AI workloads and simplify IT operations assuring exceptional secure user and application experiences—from the edge to the data center, to the cloud. Additional information can be found at www.juniper.net, X, LinkedIn, and Facebook.

Juniper Networks

Take the next step

Begin your journey to the Self-Driving Network.

Connect with an expert →

Dive deeper into the five-stage journey to the Self-Driving Network.

Explore the journey →

See how we help unlock new growth for Juniper customers.

Read case studies →

Stay informed of upcoming events and watch past ones on demand.

AI-Native NOW series →

www.juniper.net

Juniper Networks Inc.
1133 Innovation Way
Sunnyvale, CA 94089

2000839-001 EN Sept 2025