




# Under the hood of HPE's self-driving network

The AI, data, and automation engine powering next-gen networking



## Welcome to the network that pulls its own weight

IT teams are under pressure from every direction. More devices. More applications. More expectations. At the same time, resources are tighter and timelines are shrinking. Traditional networks cannot keep up. They react after problems happen and depend on manual effort to fix them.

HPE changes that model entirely, shifting IT's role from managing infrastructure to delivering exceptional user experiences. But to understand how, you have to look under the hood.

Just like a high-performance vehicle, a self-driving network is powered by a set of tightly integrated systems working in sync. These engines continuously process data, apply intelligence, and take action in real time. The result is a network that detects, diagnoses, and resolves issues before users are impacted.

With a decade-long journey toward the self-driving network, HPE is advancing that vision with agentic AI, accelerating the shift from assisted-driving to fully autonomous self-driving operations.

And this is not a future concept. It's powered today by solutions like HPE Aruba Central, Mist AI, and Marvis AI.

## Rethinking what "good enough" networks deliver

Before exploring what sets a self-driving network apart, it helps to consider how most networks operate today.

Many networks have evolved over time, bringing together multiple tools, dashboards, and processes. IT teams monitor devices, set thresholds, and respond to alerts as they arise. When issues occur, different teams often collaborate, each focusing on their area of expertise to piece together the full picture.

This approach works, but it makes it harder to quickly identify root causes or maintain consistent user experiences across the environment. In some cases, a network may appear healthy from a technical perspective while users experience performance issues at the application level.

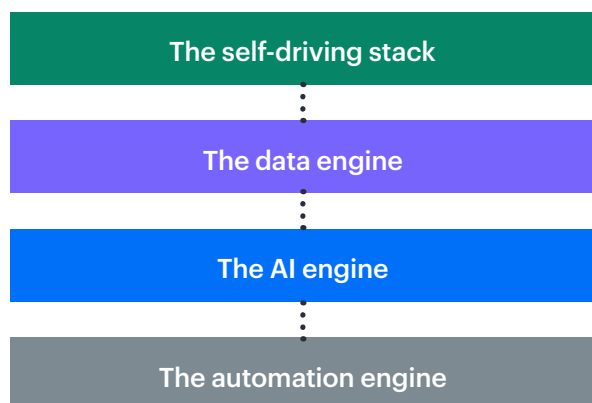
HPE takes a different approach by focusing on Service Level Experience (SLE) telemetry. Rather than only asking if the network is up, it measures whether users can actually do their work.

This shift helps enable a more coordinated system, powered by three engines working together to continuously optimize performance and experience.

## Three engines. One intelligent system.

Now let's look under the hood.

What powers a self-driving network is not a single feature, but a set of tightly integrated engines working together in real time. Each engine plays a distinct role, but none operate in isolation. They continuously share data, apply intelligence, and take action as part of a unified system.



Together, these engines create a self-driving network that can self-monitor, self-configure, and self-heal once trust is established.

## The data engine

### Infrastructure, telemetry, and data collection

The data engine is what fuels the system, while telemetry is how that fuel is generated, captured, and delivered in real time.

Across the network, telemetry is continuously streamed across wireless, wired, WAN, and data center environments, from access points and switches to client devices and applications. This includes device health, traffic flows, RF conditions, and application performance. Solutions like HPE Aruba Networking Access Points, HPE Juniper Access Points, HPE Aruba Networking CX Switch Series, and HPE Juniper EX Series Switches provide the foundation for this deep visibility.

All of this information creates a continuous stream of network, user, and application experience data. But data on its own does not provide answers.

HPE transforms this data into Service Level Experience (SLE) insights, creating a real-time view of user and application performance. Instead of relying on device status, the platform evaluates user and application SLEs against established baselines to determine whether applications are performing and users are experiencing high-quality services.

This telemetry feeds one of the largest networking data lakes in the world, built on billions of devices and real-world interactions at massive scale and it's this foundation that delivers the precision required for true self-driving operations.

### Why SLE changes the game

SLE focuses on user experience, not just device status, answering the questions that matter most:

- Are users experiencing slow applications?
- Where is the impact happening?
- How many users are affected?

### AI-native platforms: Built for AI from the ground up

All of this data is aggregated and processed within HPE Aruba Central and Mist AI.

Both platforms are built on cloud-native microservices architecture that enables a "develop once, deploy twice" mode, allowing innovations to be shared across them. New capabilities can be delivered quickly, and existing services can evolve independently, making each platform more intelligent over time.

This architecture enables the platforms to ingest massive volumes of telemetry, run AI services in real time, and orchestrate policies across the network. It also ensures innovation is continuous, not tied to rigid upgrade cycles.

### Why microservices matter

Microservices act like modular building blocks:

- Independent services evolve without disrupting the system
- New AI capabilities can be added quickly
- Platform scale with demand



## The AI engine

### AIOps intelligence

If the data engine fuels the system, the AI engine makes sense of it, turning telemetry into actionable intelligence.

By analyzing data across wired, wireless, WAN, and application domains, HPE applies AIOps to detect anomalies, predict issues, and perform automated root cause analysis. This cross-domain correlation is a key differentiator, uncovering issues traditional tools often miss.

Instead of simply reporting what happened, the platform identifies why it happened and what should happen next. This marks the shift from observability to action.

### Agentic AI framework

At the core of this intelligence is an agentic AI framework made up of specialized, collaborative agents.

Each agent operates as a domain expert, focused on areas such as wireless performance, switching, or application experience. These agents dynamically form workflows, working together to troubleshoot issues, validate configurations, optimize performance, and determine the best course of action.

This coordinated model enables the network to solve complex, multi-domain problems quickly and accurately.

### Experience models and digital twins

HPE extends AI beyond infrastructure into user experience itself.

The Large Experience Model (LEM) sits at the heart of HPE's user experience framework, analyzing billions of real-world data points from applications like Teams and Zoom to understand and predict application and user experience. In parallel, Marvis Minis simulate user sessions, continuously testing network conditions and identifying potential issues before they impact users.

Together, they shift operations from reactive troubleshooting to proactive assurance.

### AI assistants and copilots

HPE brings AI directly to IT teams through conversational interfaces.

With Marvis AI and the networking copilot in HPE Aruba Central, teams can ask questions in natural language and receive clear, actionable answers. Through copilots, the network performs root cause analysis, identifies impacted users and devices, and recommends or executes self-driving actions when permitted.

This saves operations teams significant time spent identifying, troubleshooting, and resolving issues.

### From insight to answers in seconds

Ask a simple question and get a complete answer.

- What is the root cause of the issue?
- Which users and devices are impacted?
- What action should be taken next

## The automation engine

### Autonomous actions

This is where intelligence turns into action.

The automation engine translates AI-driven insights into real-time remediation, enabling autonomous, self-driving actions. Guided by SLE insights and human-approved policies, the network can optimize performance, correct misconfigurations, and improve application experience to maintain consistently high user experience.

This is the shift from managing fire drills to predicting and fixing issues before they become customer impacting problems.

### Human-centered design

Autonomy must be trusted to be effective.

HPE takes a human-centered approach that keeps IT teams in control at every stage. Through the self-driving actions dashboard, teams can review recommendations, approve actions, and monitor outcomes before enabling full automation.

Every action is transparent and explainable, so teams understand what the system is doing and why. As confidence grows, organizations can gradually increase automation, moving from assisted operations to fully autonomous networking.

This ensures automation amplifies the impact of IT teams, empowering them to focus on higher-value outcomes.

### Automation you can trust

Built with a human-in-the-loop approach:

- Review and approve actions before execution
- Understand every decision with full transparency
- Increase automation at your own pace

## The future of self-driving networks now

### A shift that is already underway

Networking is entering a new era. As environments grow more complex and user expectations continue to rise, reactive operations are no longer sustainable.

AI is transforming how networks operate. Systems can now continuously analyze data, identify issues, and take action in real time. This reduces operational overhead, improves performance, and allows IT teams to focus on delivering better user experiences.

This shift is not theoretical. It is already happening.

### From operations to outcomes


A self-driving network changes the equation entirely. Issues are detected before users are impacted. Root causes are identified automatically. Remediation happens in real time, either with approval or autonomously.

The network continuously learns and improves with every action.

With solutions like HPE Aruba Central, Mist AI, and Marvis AI, organizations can begin this journey today. They can start with visibility, build trust in AI-driven recommendations, and gradually move toward fully autonomous operations at their own pace.

## Built for what comes next

HPE's self-driving network is built on more than a decade of real-world AI, trained on billions of devices and interactions. Unlike approaches that layer AI onto legacy systems, HPE uses a unified, AI-native architecture that delivers intelligence across wired, wireless, WAN, and cloud. By correlating data across the entire environment, the platform can identify and resolve issues end to end, while maintaining a human-centered model with full visibility and control.




## What sets HPE apart

- A decade of AI innovation and real-world learning
- AI-native platforms built for continuous evolution
- True cross-domain intelligence across the entire network
- Transparent, human-centered automation that builds trust

The future of networking is not just automated. It is self-driving. And it is already here.

## Learn more at

[HPE.com/networking](https://hpe.com/networking)



Visit [HPE.com](https://hpe.com)

[Chat now](#)

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

a00158164ENW

HEWLETT PACKARD ENTERPRISE

[hpe.com](https://hpe.com)

