Kocho Enterprise Wireless Networking Services

→ HPE Mist AI Service Definition Document

Kocho
BECOME GREATER

Executive Summary

- HPE Mist AI is a cloud-based service that uses artificial intelligence and machine learning to optimise and automate enterprise networks across wireless LAN, wired LAN, and SD-WAN domains.

- It delivers a “self-driving” network experience: the platform continuously learns from network data, identifies and resolves issues in real time, and provides deep visibility into user experiences. Mist AI is offered as a subscription service, combining cloud-managed network infrastructure with AI-powered operations and support.

Service Features & Capabilities

- **AI-Driven Network Operations:** Mist AI's cloud engine applies data science to telemetry from HPE access points, switches, and gateways, enabling automated event correlation, rapid root cause identification, and proactive anomaly detection across the network. In effect, the system can find and fix many network issues automatically (HPE refers to this as Self-Driving Network™ operations). For example, Mist uses supervised machine learning to correlate events for fast troubleshooting, and time-series anomaly detection to spot emerging problems before users notice.

- **Marvis Virtual Network Assistant:** A standout feature is Marvis, Mist's built-in virtual network assistant – the industry's first AI-driven assistant for IT networking. Marvis provides conversational troubleshooting and insight via a natural language interface. IT staff can ask Marvis questions (in plain English) about network health or client issues, and Marvis will answer with observations or recommended fixes. It continuously learns from network data and support tickets to improve its guidance over time. This AI assistant helps reduce mean time to repair and offloads routine support tasks by providing instant answers and automated actions.

- **Wireless Assurance and Optimisation:** Mist AI was initially known for revolutionising Wi-Fi management as the first AI-driven WLAN. It makes Wi-Fi more predictable, reliable, and measurable by monitoring service levels and optimising RF parameters in real time. The Mist platform sets and tracks Service Level Expectations (SLEs) – e.g. metrics like throughput, latency, and roaming time – and will automatically adjust radio settings or recommend changes to meet those targets. An AI-driven Radio Resource Management feature continuously adapts channel and power settings based on conditions, and can instantaneously react to interference or coverage issues. The result is a more resilient wireless network with less manual tuning.

- **Wired and WAN Assurance:** Mist AI extends the same AI-driven insight and automation to HPE EX Series Ethernet switches and SRX/SSR WAN gateways. This unified approach means the platform correlates wireless and wired events together (e.g. identifying if a user's issue is actually switch-port or VLAN related rather than the Wi-Fi). It also covers SD-WAN anomalies for organisations using HPE Mist for WAN assurance. All network domains – campus Wi-Fi, branch switching, and WAN – feed into one cloud AI engine, giving IT a holistic view of user experience from client to cloud. By assuring both wired and wireless, Mist helps ensure that a problem in one segment (like an upstream DHCP issue on the wired network) is identified as the root cause of symptoms observed on another segment (like Wi-Fi clients failing to get IP addresses), thus speeding up resolution.

Service Features & Capabilities Continued

- **Location Services (IoT and BLE):** In addition to network connectivity management, the Mist platform integrates location-based services using built-in Bluetooth LE radios in the access points. Mist's patented virtual Bluetooth LE (vBLE) technology enables use cases like wayfinding, proximity notifications, occupancy analytics, and asset tracking without the need for battery-powered beacon hardware. The platform uses unsupervised machine learning with these BLE signals to accurately locate devices or tags indoors. This capability allows organisations to leverage their Wi-Fi infrastructure for IoT and location applications (for example, finding a high-value asset in a hospital, or pushing a greeting to a shopper's phone when they enter a store section). All of this is managed through the same Mist cloud interface, making location an integrated service rather than a bolt-on.

- **Open APIs and Programmability:** The Mist AI cloud is fully programmable, exposing 100% of its functions via open APIs for integration and automation. This allows customers or partners to tie Mist into external IT systems – for instance, exporting data to a data lake, integrating with ServiceNow for automated ticket creation, or triggering network changes in response to events. The platform's microservices architecture and API-first design mean it can easily fit into DevOps workflows or be controlled programmatically. In practice, many tasks that traditionally required manual CLI or separate tools (config changes, firmware updates, user onboarding workflows) can be automated with scripts or third-party software tapping into Mist's APIs. This openness is a key part of Mist's "network as a service" model, enabling easy extensibility and custom integrations (e.g. using AIOps data from Mist in other IT dashboards or linking Mist with security systems for network access control).

- **Security and Zero Trust Integration:** Mist AI is designed with an enterprise security mindset. The cloud service is hosted in Tier-1, ISO 27001-certified data centres and runs on hardened Linux systems within AWS and GCP, with strong encryption for data in transit and at rest. Only necessary network ports are open to the cloud, and regular security testing (including penetration tests and vulnerability scans) is conducted on the service. In terms of network security features, Mist-based networks support standard protections like WPA3/802.1X authentication, role-based access control, IPSec tunnels, and rogue AP detection. Mist also aligns with zero-trust principles by integrating with identity and policy systems – for example, it can enforce granular access policies via HPE's cloud NAC solution and coordinate with firewall policies (HPE's Connected Security framework). The Mist cloud itself employs a "no trust" architecture for management access (all admin access is over secure channels with multi-factor authentication, etc.). Overall, customers can expect a secure-by-design service both in how the cloud is operated and in the security features available for their wireless/wired environment.

Service Delivery Model

- **Cloud-Native Architecture:** Mist AI is delivered entirely from the cloud (there are no on-premises network controllers – the intelligence resides in HPE's Mist cloud service). It utilises a modern microservices architecture deployed across multiple public cloud providers and availability zones for resiliency. This design gives the service high scalability and fault tolerance: individual microservice components can scale elastically based on load, and a failure in one component (for example, the analytics engine) will not bring down the whole service. New features and bug fixes are rolled out on a near-weekly basis to the cloud, with no downtime or maintenance window needed thanks to this microservice, continuous-delivery approach. HPE pushes updates in a rolling fashion and can innovate faster by decoupling services, meaning customers always have access to the latest capabilities without performing manual upgrades.

- **Always-On, Redundant Cloud Service:** The Mist AI cloud is architected for 24x7 availability, and HPE's service description commits to a robust SLA. The official uptime target is 99.99% availability for the Mist cloud service, measured monthly. To achieve this, Mist is deployed in geographically distributed data centres and uses clustering across regions. In the rare event of a whole cloud outage, HPE notes that on-site Mist devices (APs, switches) will continue to operate using the last known configuration – so Wi-Fi and network connectivity stay up even if cloud connectivity is lost. This offline resilience ensures that critical connectivity isn't dependent on constant cloud reachability: users already connected can remain connected and new clients can still associate using cached configurations, minimizing disruption. Planned maintenance is very minimal; HPE indicates that scheduled downtime (for major upgrades, etc.) is typically less than a few minutes and occurs only 1–2 times a year, with advance notice. In practice, most updates happen without any noticeable impact due to the microservices approach.

- **Subscription-Based Licensing:** Customers consume Mist AI via subscriptions tied to each network device or service. For example, a common package is the Mist Wi-Fi Assurance subscription for each access point, which entitles the AP to connect to the Mist cloud and provides the AI features and support for that AP. Similar subscriptions exist for Wired Assurance (per switch), WAN Assurance (per gateway), plus optional add-ons like User Engagement or Asset Visibility for location services, and Premium Analytics for extended data retention and custom reports. This modular licensing lets organisations enable only the services they need. All subscriptions include access to the Mist cloud platform (via web UI or mobile app) and its updates. The cloud service is delivered to the customer "as a Service" – meaning HPE hosts and manages all the server infrastructure. Customers simply deploy the Mist-compatible hardware on site (or can use existing supported HPE hardware), and then claim those devices in their cloud tenant. From that point, configuration and monitoring is done through the cloud. There is also an optional Mist Edge appliance for certain use cases (e.g. bridging to legacy systems or tunnel traffic on-premises), but in general no on-site controllers are required – which simplifies deployment and scaling.

Support and Service Levels

- **Technical Support:** HPE and/or Kocho provides comprehensive technical support for Mist AI as part of the subscription service. End users (or resellers managing on behalf of customers) can access HPE's Technical Support Engineers (TSE) team 24x7x365 via the Mist portal or email. This team assists with troubleshooting, configuration guidance, and problem resolution for the cloud service and Mist hardware. The support model is designed to be very responsive: upon a service request, a TSE begins diagnostics and will periodically update the customer on status, escalating internally as needed if an issue is not resolved promptly. Mist's cloud design actually aids support – because the system is cloud-managed, HPE support can often remotely view live metrics, logs, and even perform packet captures via the Mist cloud to help identify issues quickly. If a hardware component (like an AP or switch) is suspected to be faulty, the support team will initiate an RMA (Return Material Authorization) process to ship a replacement unit. Hardware replacement timelines depend on the support level purchased (for instance, some customers opt for next-business-day replacement, whereas others might have same-day "Advanced Replacement" contracts, especially for critical sites). The standard Mist subscription includes software updates and support; for hardware, HPE offers a limited lifetime warranty on many AP models and requires an active subscription for ongoing cloud service and advanced support.

- **Service Onboarding and Training:** As part of the service, HPE and/or Kocho provides extensive online resources and optional training. The Mist cloud interface includes guided setup tools and Marvis can function as a "virtual support assistant" to answer how-to questions. Additionally, customers have access to the Mist Knowledge Base and documentation library for self-service support. HPE also runs a community forum and an "ideas" portal where users can request new features. While not explicitly part of the service definition, many Mist customers work with integration partners (like HPE or certified resellers) to design and deploy their network – those partners may offer additional managed services around Mist. The Mist service description notes that if Mist AI is purchased through a managed service provider (MSP) as part of a larger offering, the support flow might differ (the MSP may provide first-line support). However, the Mist cloud and HPE backend support are always available to ensure the technical issues are resolved.

- **Performance and Reliability Guarantees:** The Mist AI cloud's SLA promises 99.99% uptime, reflecting a high level of reliability for a mission-critical network service. HPE commits to using commercially reasonable efforts to meet this target, and provides a public status dashboard (status.mist.com) for transparency during any incidents. In case of an outage or severe service issue, HPE's Cloud ops and CIRT (Cyber Incident Response Team) will respond and later provide an incident report to customers detailing the root cause and any remedial actions. This robust SLA, along with the design that even a cloud interruption won't knock users offline, gives confidence for enterprises offloading their network control to the Mist cloud. From a support perspective, critical issues (P1 cases) trigger automatic escalation to senior support management if not resolved quickly, and there are well-defined processes for bug fixes and feature requests (including regular software releases that address known issues). All of these points illustrate a mature service model backing the Mist AI platform.