NEWSLETTER

HPE Networking at Mobility Field Day

By **Keith Parsons, CWNE #3** · 2026-05-18 · originally on LinkedIn

What the Self-Driving Story Actually Means for Wi-Fi Engineers

After 90 minutes with Bob Friday, Sudheer Matta, and Miles Davis on the "*Journey to Self-Driving, Together*" session, the verdict is this: the Aruba + Mist combination is further along on a real, measurable AI story than I expected to see this soon, and most of the actual value for a working Wi-Fi engineer lives in three pieces, not in the agentic AI buzz around them.

Here's what I'd tell a junior engineer to pay attention to.

1. The Large Experience Model (LEM) is Real, and Measurable

LEM is trained on millions of Zoom and Teams sessions joined to network telemetry. It predicts video collaboration experience for every user on the network every minute, and uses Shapley math to surface which network feature (latency, packet loss, capacity, jitter) is driving the bad minutes.

That's not "AI on a slide." That's a model that points at the real RF problem behind a bad Zoom call. Bob's claim, *every single customer who deployed LEM saw a material improvement on Teams and Zoom experience,* is the kind of claim you verify with one Proof of Concept.

For a working Wi-Fi engineer, this is the thing. Most of us still get tickets that read "Zoom is bad." Now there's a tool that closes the gap between "Zoom is bad" and "your 5GHz airtime is saturated at this AP between 2 and 4 PM." That's worth knowing.

2. Marvis Minis Are Synthetic Digital Users, Now Deployed in Global Cloud POPs

Marvis Minis run synthetic transactions from the AP, switch, or router, and now from cloud POPs around the world. They measure first-hop and end-to-end behavior, packet-capture the failures automatically, and tell you whether the problem is on the LAN, WAN, or app side.

The reason this matters for a Wi-Fi engineer: most of us spend the early part of a ticket trying to prove the problem is NOT Wi-Fi. Marvis Minis answers that question before the ticket even hits your queue.

Honest open question, and credit to whoever asked it: what happens to UXI sensors? The honest answer from the panel: nothing yet. UXI is staying, Minis is staying, the future is convergence on the richest agent. If you've already deployed UXI, you didn't waste your money. If you haven't, hold and watch.

3. The Trust List is the Honest Part of the Self-Driving Story

This is the part that should make every operator pay attention. Marvis can self-drive certain actions, but only the actions IT explicitly puts on the trust list. Everything else still needs a human approval.

Demo example: a rogue DHCP server gets detected on a wired port, and Marvis can push the remediation config automatically, but only if you've turned that specific action on. You can roll it back from the audit trail.

This is the right design. You don't get into a Waymo until it's driven a million miles without an accident. The team said the quiet part out loud: "*Efficacy is the name of the game when it comes to agentic*." Their internal benchmark, 70 to 80% Marvis resolution rate on real customer support tickets, measured over years, is the kind of number you want to see before you let any AI touch your config. Push that to 90%, and the trust list grows.

Though I do understand why getting the last 10% is going to be increasingly difficult. We've already picked off the low hanging fruit, now the final 10% of issues are quite complex and not as easy to solve.

Walk in With Eyes Open

A few things didn't get fully answered, and a working engineer should know:

- **Naming is a mess right now.** Marvis, Copilot (not Microsoft Copilot), agentic mesh, Aruba Central, Mist, the delegate panel hoped the branding gets resolved at HPE

Discover Las Vegas. Until then, the same engine wears two badges depending on which dashboard you log into.

- **The Aruba + Mist convergence is engineering-real, not GA-real.** The Marvis data science team is one team. The agentic framework is one framework. The customer-facing experience is still two clouds. Expect the seam for a while longer, perhaps a long while longer.
- **Roaming is still mostly an RRM problem, not an algorithm problem.** Credit to Sudheer for saying it on stage: a lot of band-steering and load-balancing got phased out because Apple figured out how to work around it. The fix for bad roaming is good design, coverage, channel plan, power. The agent helps you find the design problem faster. It can't fix the coverage hole you didn't design for.
- **Determinism in the LLM answers is still being worked.** Sam Clements asked the right question, if two operators in the same org ask the same question, do they get the same answer? Today: similar, not identical. The team is building episodic / long-term memory to fix it. Worth tracking.

The Verdict

The "*Journey to Self-Driving, Together*" session was the strongest "*we did the hard work*" presentation I've seen from a major Wi-Fi vendor in this cycle. The data is real, the model has efficacy numbers behind it, and the trust list is the right operator-friendly answer to "*should I let AI change my config*." Nothing here was vapor.

For a WLAN engineer trying to level up: focus on the LEM, learn how to read its outputs, and stop spending the first half of every ticket trying to prove Wi-Fi is or isn't the problem when a Mini can answer that for you in seconds.

If you're running Aruba today, this story is coming to your dashboard. If you're running Mist today, you're already in it. If you're running neither, watch this space and let the convergence settle before you pick a side.

Mobility Field Day session videos are posted publicly, YouTube may be the easiest to track down. Watch it yourself before you let me convince you.

#WLPC #WiFi #MobilityFieldDay #HPENetworking #Aruba #JuniperMist #MFD14

By Keith Parsons, CWNE #3, Wireless LAN Professionals. Free to read and share. Facts