How AI Makes Better Network Engineers

Tom Hildebrand | Ubiquiti
tom.hildebrand@ui.com | WLAN Pros Slack | LinkedIn

Let’s talk live events

Live event networks are some of the most challenging.

No chance to evaluate a “baseline” or benchmark

Limited setup time can lead to mistakes

Hard to review data during the event

Multi-engineer, multi-company teams inevitably lead to friction.

Event Stats

Two buildings, 300+ APs, 500+ switches, network up and running for a week.

17k wireless devices, the same or greater number of wired clients for 8500 people.

Laptops, not just phones, 24H/day

100+ terabytes of data usage forecast



The Venue on Day 0





CENTRE - MAIN FLOOR

Managing a NOC

The larger the event, the more information overload you have to deal with.

Every AP, switch, VLAN, and gateway create an entire dashboard's worth of info.

Dashboards are for observability, but not actionability!



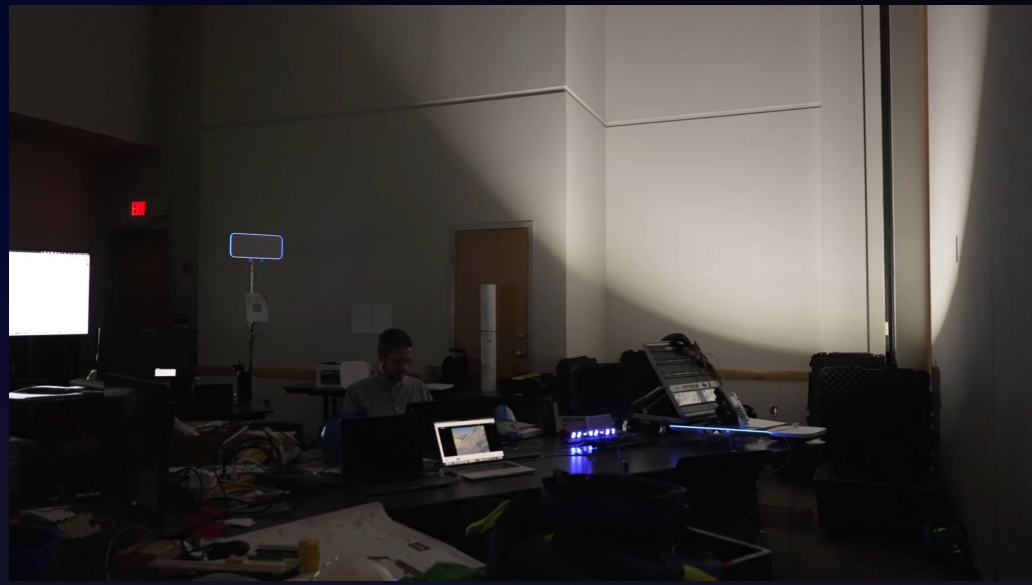
Managing a NOC in 2025

Last year we had 10-15 people manually monitoring different parts of the event.

Manually comparing notes and validating user reports took a lot of time.

Central network management didn’t mean central human visibility or quick decisions.

Though the event was successful, it was borderline unmanageable… for humans.

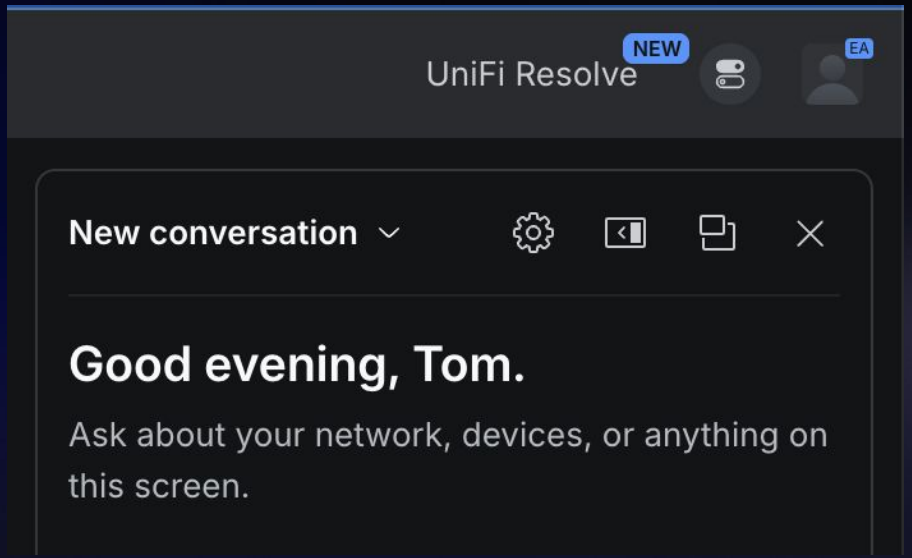


AI was transformative in 2026

On the customer side and at Ubiquiti, we both deployed AI pilot projects to make the NOC more efficient.

AI was a force multiplier for our staff, not a staff replacement. Let it tackle the small tasks while you focus on the bigger picture.

With our AI helpers, including UniFi Resolve, our teams were able to make more intelligent decisions, more quickly, with less stress.

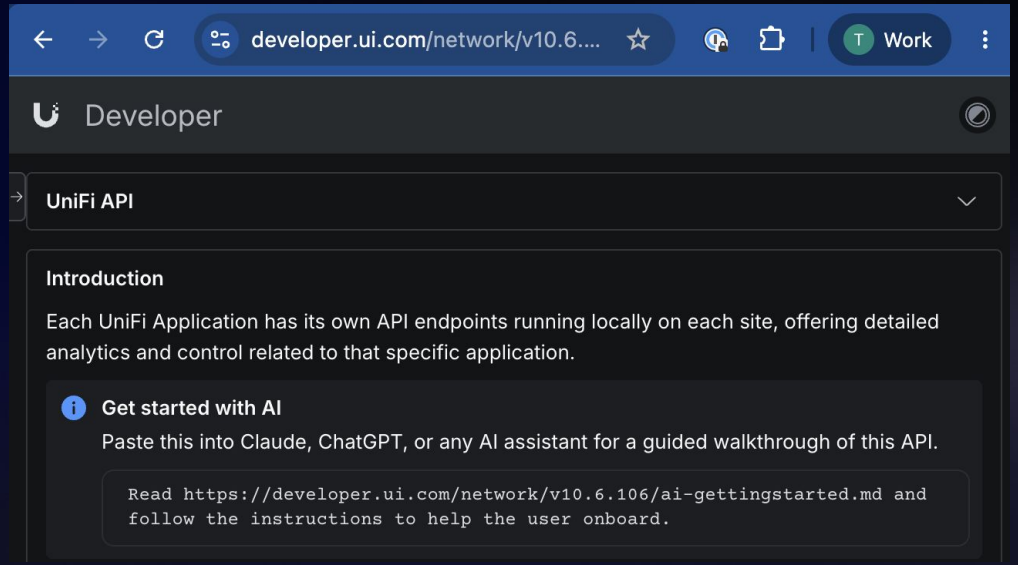


What's in an AI?

AI models continue to improve, but tooling and data access are key.

Using an AI on a production network requires:

- Accessible APIs and simple data schema.
- Well organized technical documentation.
- A single source of truth for statistics.
- Support for tool calls.

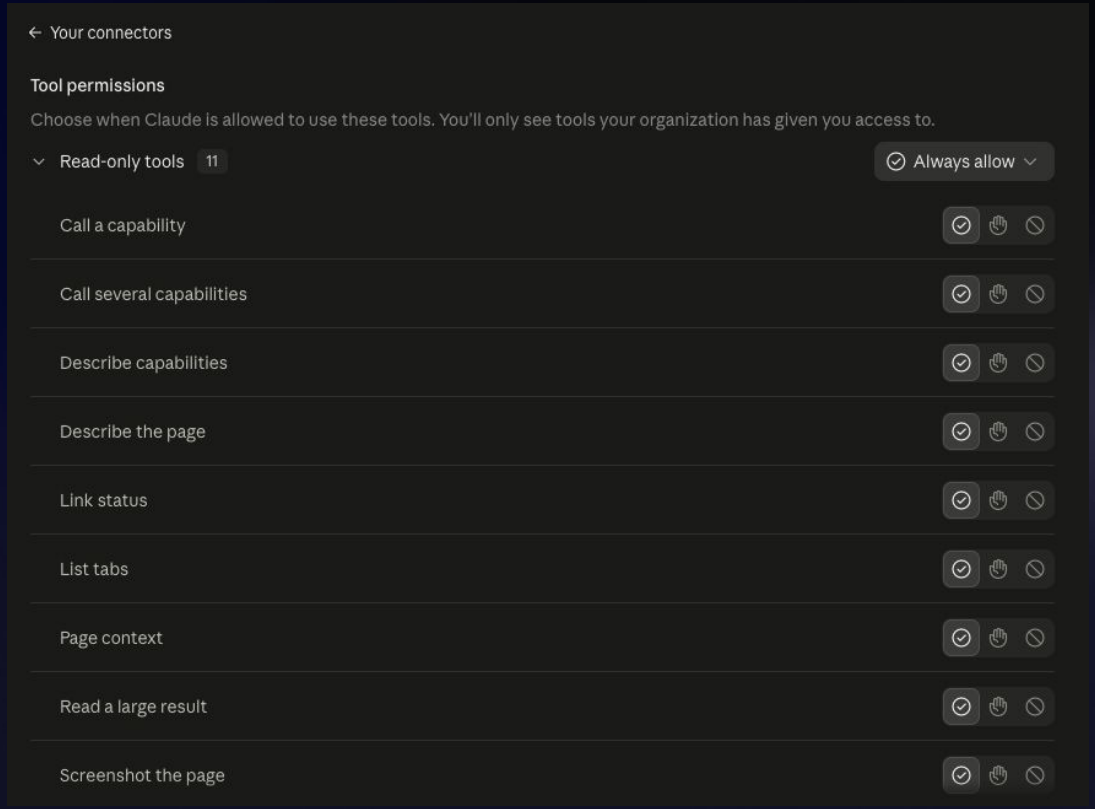


A Brief Note on Permissions and APIs

Both AI tools were scoped with read-only access to the network.

AI tools are very powerful, but still hallucinations and unexpected behavior can happen. The easiest safeguard is permissions management.

UniFi Resolve is currently read only. While we have APIs that allow for write access, we do not want to include them by default.



Teaching your AI your equipment

When intaking an API directly, KBs are great, but an OpenAPI JSON is even better.

Models have gotten a lot better at scanning the web for KB articles - use Generative Engine Optimization (GEO) when writing!

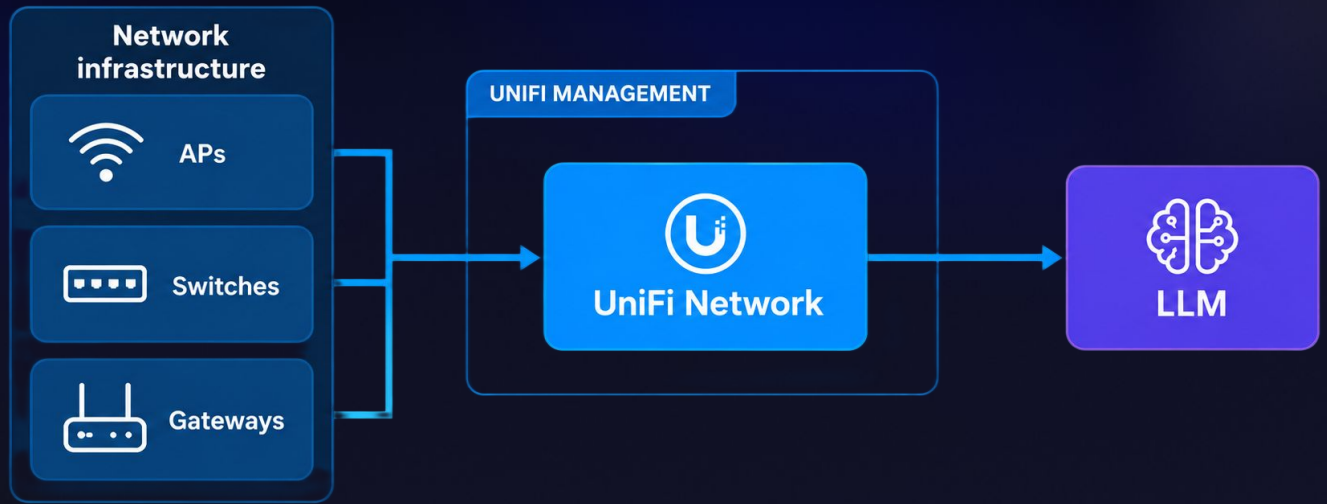

```
{
  "components": {
    "schemas": {
      "ACL rule": {
        "discriminator": {
          "mapping": {
            "IPV4": "#/components/schemas/IntegrationIpAclRuleDto",
            "MAC": "#/components/schemas/IntegrationMacAclRuleDto"
          },
          "propertyName": "type"
        },
        "properties": {
          "action": {
            "description": "ACL rule action",
            "enum": [
              "ALLOW",
              "BLOCK"
            ],
            "example": "ALLOW|BLOCK",
            "type": "string"
          },
          "description": {
            "description": "ACL rule description",
            "type": "string"
          },
          "destinationFilter": {
            "description": "Traffic destination filter"
          },
          "enabled": {
            "example": true,
            "type": "boolean"
          },
          "enforcingDeviceFilter": {
            "$ref": "#/components/schemas/ACL rule device filter",
            "description": "IDs of the Switch-capable devices used to enforce the ACL rule. When null, the rule
will be provisioned to all switches on the site."
          },
          "id": {
            "format": "uuid",
            "type": "string"
          },
          "index": {
```

Where's all the data?

We have a centralized database of historical and current network statistics, including data from each AP, switch, and gateway under management via UniFi Network

UniFi Resolve uses an MCP server to talk directly to the UniFi Network database and APIs via an open browser window's connection.

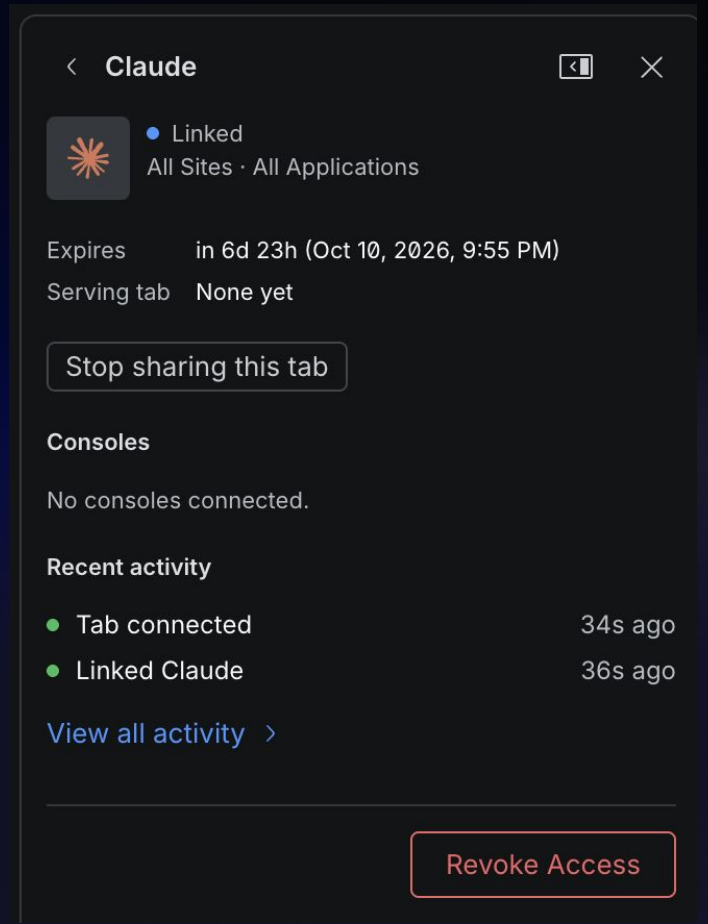
Your data remains private and yours - plug it into your LLM of choice.



Data Privacy and Audit Logging

Connected models are logged, including the actions/APIs called and what sites were accessed.

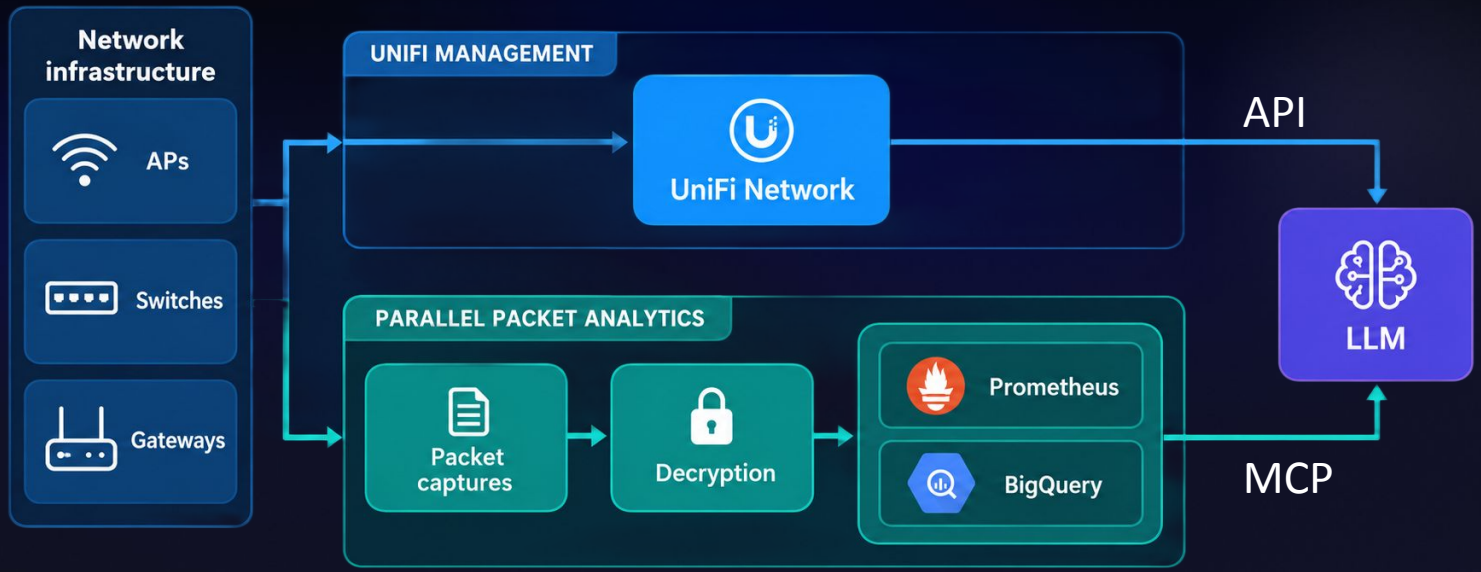
When using any API, even for read only actions, keep an audit log!



Where's all the data?

Our customer built a very similar system, but without native MCP calls to Network.

This required self-hosted infrastructure, but accomplished the same goal.

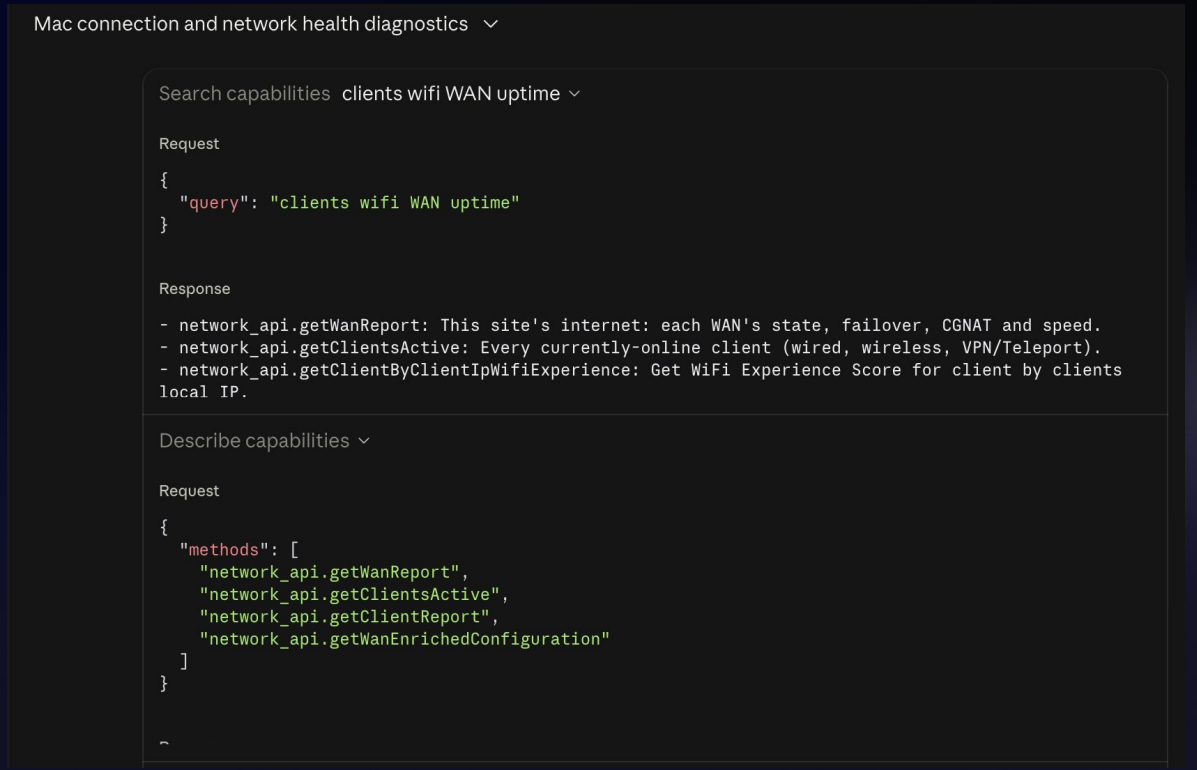


Giving your model tool calls

Tool calls take a model from general networking knowledge to knowing your specific environment.

Let the LLM process your language, then use tools to interpret the network's data.

Manufacturers need to build these translation layers for the models to work their best.



Using AI Live - RF Checks

A common question during these events is RF conditions. Every vendor has a dashboard for RF stats, but drilling down into each AP and comparing can take significant time.

A quick ask from ChatGPT with our MCP server gives us this response.

We immediately knew where the poorly performing APs were, which channels were concerning, and the model also flagged high TX retries - something not directly asked but clearly relevant.

Focusing only on online APs and 5/6 GHz radios (183 online APs, 366 radios checked):
Overall 5/6 GHz load is healthy on average

- 5 GHz average channel utilization: 58%
- 6 GHz average channel utilization: 58% Those averages are fine, not saturated.

But 15 radios are running hot (80%+ utilization) — a mix of both bands, not one:

- `C3T-C-149-149-ACC2` (5 GHz, ch149): 99%
- `C3T-S-149-149-B0B2` (5 GHz, ch149): 98%
- `A1CS-NE-132-133-5052` (5 GHz, ch132): 92%
- `B3CS-SW-136-101-732C` (5 GHz, ch136): 89%, plus 42k tx retries
- `Gaming-E-E-36-165-3BA6` (5 GHz, ch36): 86%
- `HT4-157-197-ABB4` (6 GHz, ch197): 85%
- `C3T-N-132-133-7860`, `D3CN-NW-36-165-575A` (5 GHz): 84%
- `B1CS-SE-136-101-E5E0` (6 GHz, ch101): 84%, plus 22.9k tx retries
- Several more in the 80-84% band, both 5 GHz and 6 GHz

A further 187 radios sit in the 60-79% range, so this is a moderately loaded network, not a crisis, but the top ~15 radios above are worth attention.
Retry rates are the sharper signal here — several radios have very high tx retries even where utilization isn't extreme, which usually points to RF interference or weak client links rather than pure congestion:

- `A2T-C-149-149-FEAC` (6 GHz, ch149): 149,637 retries at only 71% utilization
- `B4T-S-165-197-AB30` (6 GHz, ch197): 140,300 retries
- `B2CS-NW-100-5-7770` (5 GHz, ch100): 128,701 retries
- `B1T-S-165-197-7956` (5 GHz, ch165): 111,327 retries
- `B1CN-NE-40-21-DA64` (5 GHz, ch40): 85,333 retries

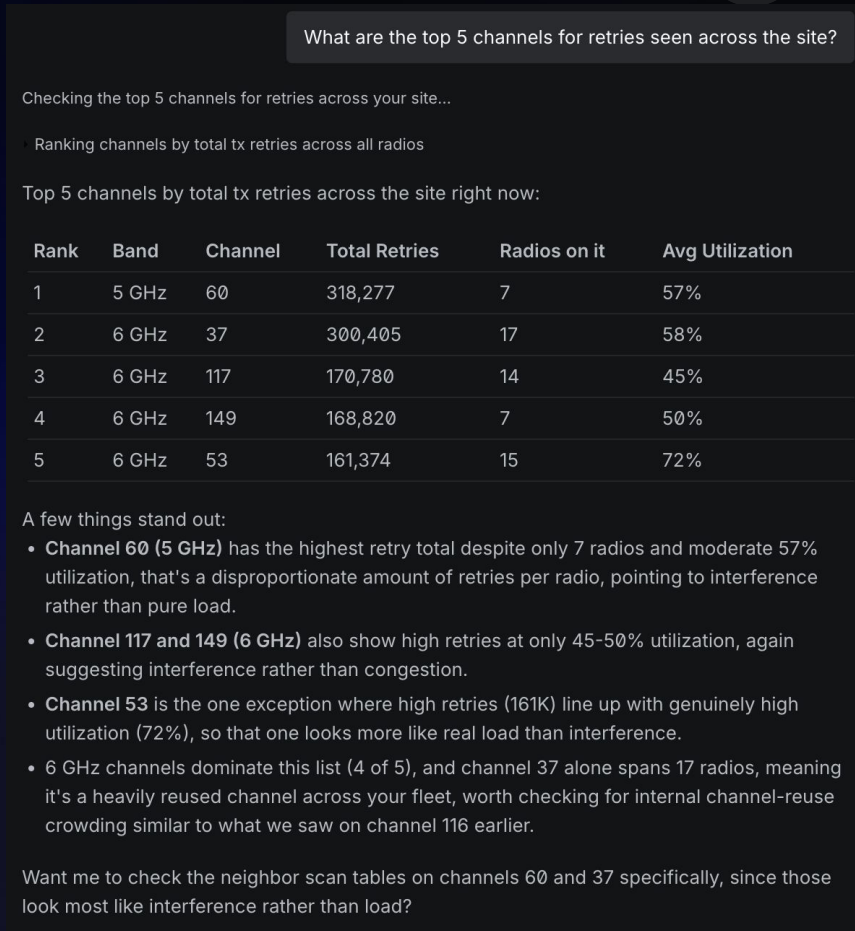
High retries with moderate utilization on channels like 149, 165, and 197 suggests co-channel interference or DFS-related retransmits on those specific 5/6 GHz channels rather than raw client overload. I don't have per-SSID client counts available to isolate Summit-WiFi traffic specifically from other SSIDs on these same radios, since the client list is too large to pull in one read on this controller, so this reflects total radio load, which is dominated by whichever SSID carries the most clients on each AP.
If you want, I can pull the channel plan for the worst offenders (149, 132, 136, 197, 100, 165) to check if they're on overlapping channels causing the interference.

Using AI Live - RF Checks

An AI, with access to the right data points, is incredible at finding signal from noise.

Its ability to draw conclusions autonomously is far less interesting than finding correlations between data sets! Here, channel interference vs number of radios not tracking as expected.

This told us to check our channel plan and look for interferers.



What are the top 5 channels for retries seen across the site?

Checking the top 5 channels for retries across your site...

Ranking channels by total tx retries across all radios

Top 5 channels by total tx retries across the site right now:

Rank	Band	Channel	Total Retries	Radios on it	Avg Utilization
1	5 GHz	60	318,277	7	57%
2	6 GHz	37	300,405	17	58%
3	6 GHz	117	170,780	14	45%
4	6 GHz	149	168,820	7	50%
5	6 GHz	53	161,374	15	72%

A few things stand out:

- **Channel 60 (5 GHz)** has the highest retry total despite only 7 radios and moderate 57% utilization, that's a disproportionate amount of retries per radio, pointing to interference rather than pure load.
- **Channel 117 and 149 (6 GHz)** also show high retries at only 45-50% utilization, again suggesting interference rather than congestion.
- **Channel 53** is the one exception where high retries (161K) line up with genuinely high utilization (72%), so that one looks more like real load than interference.
- 6 GHz channels dominate this list (4 of 5), and channel 37 alone spans 17 radios, meaning it's a heavily reused channel across your fleet, worth checking for internal channel-reuse crowding similar to what we saw on channel 116 earlier.

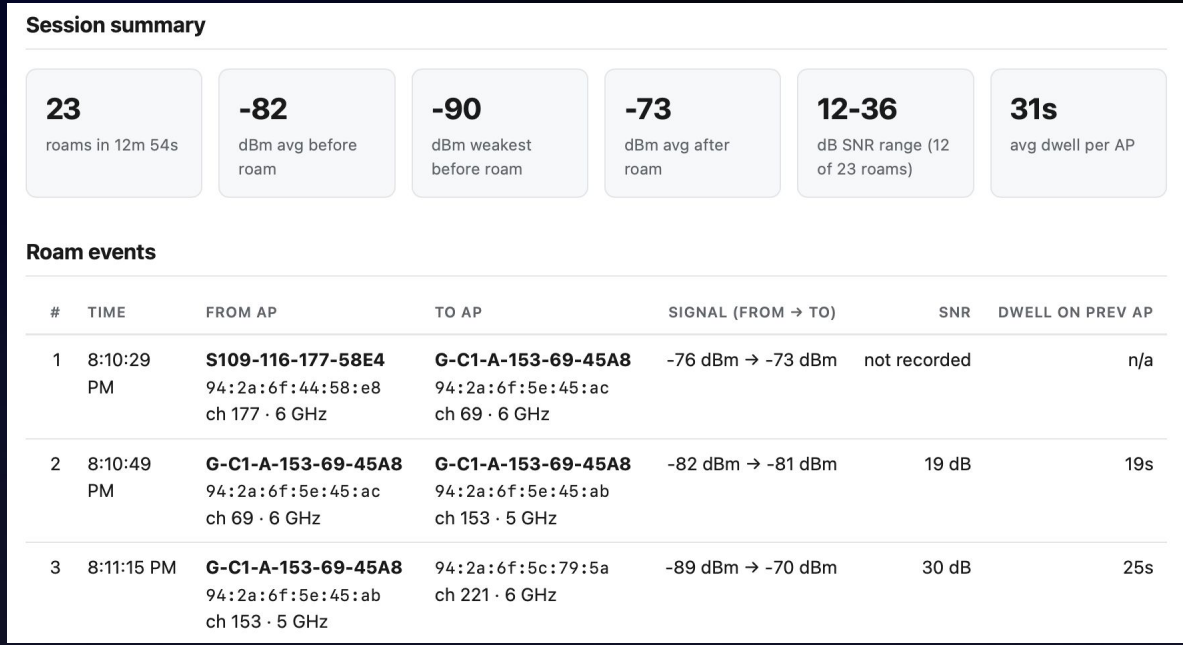
Want me to check the neighbor scan tables on channels 60 and 37 specifically, since those look most like interference rather than load?

Using AI Live - Roaming Analysis

Using the WLAN Pros Toolbox “Roaming Log” feature, we were able to reliably benchmark client roaming behavior across MacBook gens (M1 to M5).

Rather than parse the logs by hand, a perfect job for an LLM.

This helped fine tune AP positioning and channel planning, beyond what surveys can show.



Session summary

23	-82	-90	-73	12-36	31s
roams in 12m 54s	dBm avg before roam	dBm weakest before roam	dBm avg after roam	dB SNR range (12 of 23 roams)	avg dwell per AP

Roam events

#	TIME	FROM AP	TO AP	SIGNAL (FROM → TO)	SNR	DWELL ON PREV AP
1	8:10:29 PM	**S109-116-177-58E4** 94:2a:6f:44:58:e8 ch 177 · 6 GHz	**G-C1-A-153-69-45A8** 94:2a:6f:5e:45:ac ch 69 · 6 GHz	-76 dBm → -73 dBm	not recorded	n/a
2	8:10:49 PM	**G-C1-A-153-69-45A8** 94:2a:6f:5e:45:ac ch 69 · 6 GHz	**G-C1-A-153-69-45A8** 94:2a:6f:5e:45:ab ch 153 · 5 GHz	-82 dBm → -81 dBm	19 dB	19s
3	8:11:15 PM	**G-C1-A-153-69-45A8** 94:2a:6f:5e:45:ab ch 153 · 5 GHz	94:2a:6f:5c:79:5a ch 221 · 6 GHz	-89 dBm → -70 dBm	30 dB	25s

Conclusion

AI is a force multiplier for you, not your replacement.

Let AI enable you to focus on the most impactful parts of your work.

Spend less time in the NOC and more time on the floor and with your customers.

Don’t be a meat proxy! What does the model say in *your* context? In *your* deployment?