

Three new capabilities

Network-layer detection for IT, OT and AI — passive, agentless, and never in the path of your traffic.

1 AI security — seeing the AI on your wire

The gap. Endpoint tools, cloud access brokers and firewalls see a connection, not whether the thing on the other end of it was a person, an agent, or an attacker wearing an agent's name. Application logs only exist where someone thought to add them — and never for the systems nobody declared.

- Shadow AI: which AI services staff and systems actually use, and what they send out — a live inventory, usually the first surprise.
- Fake AI crawlers: identity established from the network, not the label anyone can type to get past your rules and rate limits.
- Your own AI estate: model extraction, weight exfiltration, corpus tampering, and agents reaching unvetted external tool servers.
- AI-driven reconnaissance of an OT estate, caught before anything reaches a controller.

2 AI drift and rogue AI — two failures, reported separately

The distinction. Rogue is a scope question: your approved AI reaching a provider, model or tool nobody approved, or acting unattended outside declared hours. Drift is a quality question: the same approved AI quietly getting worse because a provider rolled a model without telling you. Most vendors sell them as one thing, which is convenient and wrong.

- Rogue behaviour detected on the wire: destinations, models and tools outside your register, new external tool servers, unattended autonomy, consumption runaway.
- Drift reported as dated indicators, with an honest warm-up — nothing is asserted until fourteen daily samples are held.
- With your gateway metadata, a model or version change becomes a stated fact rather than a hypothesis, and token, truncation and refusal rates become measurable.
- We will not call a number derived from packets a quality score. Where the answer quality question needs your scoring, we say so.

3 Brand impersonation and mobile app sideload

The economics. A lookalike domain costs about twelve dollars and a valid certificate is issued within minutes. And the payload no longer arrives as a phishing page: a fake app cannot come from the official store, so it arrives by sideload — and that arrival crosses the network you already let us watch.

- Lookalike names caught at registration: public certificate logs watched every 30 minutes, scored against your brand rather than matched literally.
- Registration is reconnaissance; contact from inside your network is the attack. Both reported, only one treated as an incident.
- Sideloaded apps and configuration profiles caught on arrival — no setup, no brand word, live on day one.
- Which handset, which source, what it fetched, and whether the same payload came back renamed.

What we will not claim

We do not grade answer quality from packets. We do not detect prompt injection, jailbreaks or hallucination — those live in the prompt and tool-call layer, which no passive sensor sees. A handset on mobile data never crosses your network and is invisible to us. We do not block anything; we observe, and your inline component enforces. Negative controls — what we deliberately do not alert on — are demonstrable on request, against a packet capture you already hold.

Next step

The fastest evaluation is a scoped pilot against a packet capture you already hold — days, not quarters. Nothing is installed on an endpoint, no application changes, and you are looking at your own network rather than a vendor lab. We will show you what we do not alert on as readily as what we do.