

STRAIGHT ANSWERS

The questions that actually decide whether an evaluation goes ahead.

These are the questions we are asked most often, including the awkward ones about limitations. A vendor that only answers the flattering questions is not worth the meeting, so the limits are here alongside the capabilities.

HOW IT WORKS

Do you install anything on our devices?	No. Nothing is installed on any observed device. WireTrace connects to a SPAN port or network TAP and observes traffic that is already there.
Do you scan or probe the network?	No. Nothing is transmitted onto the monitored network. There is no active mode, so there is no setting that could enable one by mistake.
How does it identify a device with no agent?	From what the device reveals in the protocols it already speaks: how it announces itself, what it negotiates, what identifiers it presents and how it uses its own protocols. Identity is assembled across everything it says rather than from one indicator.
How long before it is useful?	Assets begin appearing and classifying within minutes of connection. Behavioural context improves as the environment goes through its normal cycles, so a longer window gives a better picture.
Will it affect network performance?	No. A mirror port or TAP produces a copy of traffic. WireTrace receives that copy and does not participate in the network it observes.

LIMITS, STATED PLAINLY

Can you decrypt traffic?	No, and we do not attempt to. A great deal is still recoverable from encrypted sessions without touching the payload: who is communicating, using which protocol version, presenting which certificate, with what characteristics. That is often the security-relevant part.
What can you not see?	Anything not mirrored to a sensor. Coverage is a direct consequence of sensor placement, which is why placement is designed deliberately rather than assumed. Traffic that never crosses an observed point is invisible, and the platform is explicit about that rather than filling the gap with inference.
Do you store packets?	Not as a matter of course. WireTrace retains understanding of devices and their behaviour, plus the observations supporting findings. It is not a recording appliance.
Is it a SIEM?	No. It produces asset intelligence, exposure and evidence, and feeds them to the systems that carry your workflow, including a SIEM if you have one.
Is it a vulnerability scanner?	No. It correlates vulnerabilities against firmware and software observed on the wire and against real reachability. It never interrogates a device to test one.
Does it replace our endpoint tooling?	No. Endpoint tooling covers devices that accept agents. WireTrace covers the part of the estate that structurally cannot. They are complementary, and we say so in front of the incumbent vendor.

DATA, PRIVACY AND SOVEREIGNTY

Does any data leave our network?	No. WireTrace is fully on premises. There is no cloud component, no telemetry to a vendor service, and no external dependency for any capability.
Does it work in an air-gapped network?	Yes, as a supported deployment model rather than a reduced mode. That includes analytics and the AI assistant, which run inside the environment.
How does the AI work without a cloud?	Rumi is deployed within your environment and reasons over your own observed data. Nothing is sent anywhere, which is why isolation costs no capability.
What about patient or personal data?	The platform retains understanding of devices and behaviour rather than the content of clinical or business traffic, and retention is configurable. The detailed position is in the security and privacy documentation provided during evaluation.
Can you administer our deployment?	No. The platform is operated by you. There is no vendor-held administrative path into a deployed system.

DEPLOYMENT AND OPERATION

What do we need to provide?	A SPAN port or TAP at the points where visibility is wanted, and somewhere to run the platform. Plus a conversation with someone who knows the network, so placement is designed rather than guessed.
How many sensors will we need?	It depends on network topology rather than device count, because a sensor sees what is mirrored to it. Sizing is worked through during evaluation against your actual diagram.
Does it work across multiple sites?	Yes. Sensors distributed across sites report to a central platform, giving one estate-wide picture including relationships that cross between sites.
Can we separate business units or customers?	Yes. One platform can present separated environments to separate audiences, which is used both by groups with independent subsidiaries and by managed service providers.
How does it integrate with what we have?	Inventory, findings and evidence are made available to your existing security and operations tooling, so the visibility improves decisions in the systems your teams already use.

EVALUATION

What does a proof of concept involve?	A representative segment, a mirror port, and an observation window. You receive what was found on your traffic: the inventory, the exposure and the evidence that could be produced.
Will operations approve it?	This is usually the shortest part of the conversation, because nothing is installed and nothing is transmitted. Operations can watch the connection being made and verify that for themselves.
What if you find nothing interesting?	That would be a genuinely useful result and we would tell you. In practice the gap between an observed inventory and a recorded one is where the first surprise usually appears.

A question that is not here

If the question that matters to you is not on this page, it is probably the interesting one. Ask it directly and we will answer it plainly, including when the answer is that WireTrace is not the right tool for what you are trying to do.

[Book a demo at wiretrace.io/request-demo](https://wiretrace.io/request-demo) · or email sales@wiretrace.io