


WIRETRACE TECHNICAL OVERVIEW

How WireTrace works, what it supports and how it fits your network

The architecture, deployment models, protocol coverage, analytics scope, data handling and platform security of WireTrace, stated precisely enough to design, review and procure against.



For: Security architects, network and OT engineers, procurement and technical evaluators

WireTrace is a **Deep Protocol & Asset Intelligence (DPAI)** platform. Sensors receive a mirrored copy of network traffic, decode it to the level of individual protocol fields, and a central platform turns that evidence into durable device identity, classification, detection, risk and policy decisions across IT, OT, IoT and medical devices (IoMT). This document describes how it is built, how it is deployed and what it does and does not do.

250+

protocols decoded,
payloads as well as
headers

600+

built-in classification rules,
plus your own

90+

detection rules mapped to
MITRE ATT&CK and
ATT&CK for ICS

Minutes

from a mirrored port to the
first identified devices

01 Platform flow

Every WireTrace capability sits on one chain of evidence, shown on the cover. Each layer consumes what the layer before it produced, so the identity that names a device in the inventory is the same identity that scopes a detection, ranks a vulnerability and travels with an alert into your SIEM.

Layer	What actually happens
1. Network communication	Mirrored traffic from SPAN, TAP or ERSPAN; optional, governed enrichment for IT and network devices.
2. Deep protocol intelligence	250+ protocols decoded to their fields, each tied to the observation that produced it.

Layer	What actually happens
3. Asset and entity identity	One durable identity per physical device, across addresses, interfaces and sensors, with a confidence.
4. Classification and context	Domain, type, role, manufacturer, model and firmware where stated, OS, Purdue level, applications.
5. Behavioural baseline	Operator-approved baselines of communication pairs and of industrial protocol fields and values.
6. Risk and threat intelligence	Detection, threat indicators, vulnerabilities, end of life, exposures and the Asset Risk Score.
7. Policy and segmentation	Traffic rules, IEC 62443 zones and conduits, violations and governed automation.
8. Integrated control	Syslog, webhooks, email, SNMP traps, exports and a read-only REST API towards your tools.

02 Reference architecture

A WireTrace deployment has four parts: **sensors** at the points where traffic is mirrored, a **central platform** on a server or virtual machine you own, an **integration layer** towards the tools you already run, and two optional paths that are off until you configure them: **active enrichment** and, for air-gapped sites, an **offline data path**.


Does WireTrace ever transmit onto the monitored network? The capture interface never does: it only receives. WireTrace is **passive by default**, and a deployment that uses only SPAN, TAP or ERSPAN sends nothing to the devices it observes. If you choose to enable **optional, administrator-governed active enrichment**, its read-only requests leave the platform from the management interface, only towards scopes you define, and an OT-safe profile keeps them away from industrial and medical devices unless you explicitly allow it. Nothing is installed on the devices being monitored in either case.

Sensors

A sensor has two network roles. The **capture interface** receives mirrored traffic and decodes it; it has no transmit path. The **management interface** carries the sensor's results to the central platform. Sensors are enrolled with an activation token, and sources that are not enrolled are refused. The platform includes a built-in sensor; additional sensors are sized to the link, from small appliances to high-throughput servers.

Central platform

The platform installs on customer-owned servers or virtual machines from one self-contained package, and its resources are sized from the host automatically. It holds the inventory, analytics, baselines, policies, console, reports, REST API and Rumi, the private AI assistant, and builds one de-duplicated inventory across all sensors.



Reference architecture. The capture interface of every sensor only receives a mirrored copy of traffic and has no transmit path onto the monitored network. Anything WireTrace sends, it sends from a management interface: sensor reports to the platform, integrations, and, only if an administrator enables it, scoped enrichment. On air-gapped sites the platform never connects out; data arrives by import.

Integration layer and optional paths

Outbound, the platform forwards alerts and decisions to your tools and exposes a read-only REST API that is off until enabled; inbound, it accepts sign-in, device syslog and SNMP traps, indicator files and read-only switch data (section 09). **Active enrichment** adds read-only switch and router polling and credentialed collection from hosts an administrator chooses. **The offline data path** refreshes threat indicators and vulnerability data on air-gapped sites: a collection script runs on any internet-connected machine, and its checksummed bundle is carried across under your own transfer process and imported in the console. The appliance never connects out.

03 Capture options

WireTrace builds the asset inventory from mirrored traffic. Discovery is only as complete as the traffic the sensors see, so placement matters more than any other deployment decision.

Method	How it works	Use it when
SPAN	A switch copies the traffic of selected ports or VLANs to a port connected to the sensor's capture interface.	The switch supports mirroring and has spare capacity. The quickest way to start, including for evaluations.

Method	How it works	Use it when
TAP	A network TAP copies a physical link in hardware and delivers the copy to the sensor.	You need a complete, independent copy of a critical link, for example between control and supervisory levels, without touching switch configuration.
ERSPAN	A switch encapsulates its mirrored traffic and sends it across a routed network to a sensor elsewhere.	A remote site, cabinet or substation has no local sensor, and the routed path has the capacity to carry the mirrored copy.

Good placement points are the boundaries where devices actually communicate: between Purdue levels, at the industrial DMZ, at core and distribution switches carrying east-west traffic, and on links to remote sites. Devices behind routers appear by address unless their own traffic is seen or enrichment adds detail.

04 Deployment models

Single site

One platform with its built-in sensor on a server or virtual machine, fed from one or more SPAN ports or TAPs. Typical for a plant, hospital campus, building estate or evaluation. The first devices appear within minutes of traffic arriving.

Multi-site, distributed sensors

Sensors at each site or zone report to one central platform over the management network. Each device keeps one identity even when several sensors see it. ERSPAN brings in traffic from switches where no local sensor is placed.

Air-gapped

Installation, analysis, Rumi, end-of-life data and identity and classification knowledge all run on the appliance with no internet access. Upgrades arrive as offline release packages; threat and vulnerability data arrive through the offline data path. See the Air-Gapped and Disconnected Operations brief.

Segregated tenants

One platform can host segregated tenants, for example sites or business units, each with its own assets, users and permissions. Suited to organisations that run several operating companies or sites from one central team.

Sizing. Server and sensor sizing by link throughput and asset count is published in the WireTrace Deployment and Sizing Guide. Network connectivity requirements for sensors, the platform and integrations are provided during solution design.

05 Protocol coverage

WireTrace decodes **250+ protocols**. It decodes payloads, not just headers: it extracts identity and operational fields such as vendor, order number, firmware, station and host names and certificate details, and records which observed evidence produced each attribute. Many industrial, building and network protocols carry the device's own identity record, and WireTrace uses it as primary evidence. Representative protocols by family:

Family	Protocols (examples)
Industrial	Modbus, Siemens S7comm and S7CommPlus (session level), EtherNet/IP and CIP, PROFINET, DNP3, IEC 60870-5-101/103/104, IEC 61850 (MMS, GOOSE, Sampled Values), OPC UA, EtherCAT, Omron FINS, Mitsubishi SLMP/MELSEC, CODESYS, HART-IP, SEL Fast Message, IEEE C37.118 synchrophasor, PRP/HSR, PTP
Building and IoT	BACnet, KNXnet/IP, Niagara Fox, MQTT, CoAP, LwM2M, ONVIF, SNMP
IT and infrastructure	DNS, DHCP, TLS, SSH, HTTP, SMB, Kerberos, LDAP, RDP, NetBIOS, mDNS, LLDP, CDP, RADIUS, syslog, NTP

Operation semantics. For industrial protocols WireTrace understands what a conversation does (reads, writes, program transfers, mode changes, exceptions), not just that two devices talked. OT command monitoring, which raises detections on those operations, covers Modbus, S7comm, DNP3, IEC 60870-5-101/104, IEC 61850 MMS, Omron FINS, Mitsubishi SLMP/MELSEC and SEL Fast Message. Other industrial protocols, including EtherNet/IP, PROFINET, OPC UA, BACnet and S7CommPlus, contribute identity and discovery.

Encrypted traffic. WireTrace does not decrypt traffic. TLS client fingerprints and certificate details are recorded for identity, hygiene and threat matching, and encrypted sessions still yield behavioural evidence. Encrypted industrial payloads are not analysed in depth.

06 Identity and classification method

What WireTrace observes

- Device-stated hardware and protocol identifiers, including industrial identity records
- DHCP, SNMP, SSH host keys and certificates
- Names from DHCP, DNS, NetBIOS, mDNS, LLDP/CDP and directory protocols
- Operating-system indicators and applications used
- Optional: switch polling and credentialed collection

What it produces

- One durable identity per physical device, across address changes, interfaces and sensors
- Domain (IT, OT, IoT, medical and network), device type and role
- Manufacturer, and model, firmware and serial number where the device states them
- Purdue level with a stated confidence and a one-line reason
- The source and freshness of every value

- 1 Collect evidence, not guesses.** Identity combines many independent kinds of evidence. An IP address locates a device but never identifies it, and directory and authentication names corroborate a host but never identify it alone.
- 2 Join on what is unique.** Evidence is joined only on values a device states uniquely. When two devices present the same supposedly unique value (a cloned image, a vendor default, a relaying gateway), WireTrace refuses to merge them blindly and records the contradiction.
- 3 Track time.** Identity evidence is time-bounded: first seen, last seen and still valid. Address changes are tracked as transitions, not conflicts, and stale bindings expire safely.

- 4 **Classify with stated confidence.** 600+ built-in classification rules and a hierarchical device taxonomy assign domain, type, role and Purdue level. Devices without enough evidence stay unknown rather than guessed. Administrators can add their own rules in a no-code editor, with a preview before anything is applied.
- 5 **Explain.** Every identity carries an evidence-weighted confidence, and every value shows its source and freshness. The full grade and explanation are available through the API. For each device, WireTrace also states which kinds of evidence are present, missing or stale, so an empty result is never presented as a clean one.

Manufacturer is normalised to one name, with a labelled low-confidence fallback to the network-card maker. Operating system is inferred passively with a stated confidence, and credentialed collection adds the exact version. 150+ cloud, SaaS, vendor-cloud and infrastructure applications are recognised, and you can add your own.

07 Detection and analytics scope

All detection is deterministic and states the evidence it rests on. The table states the scope of each area and the condition, where one applies.

Area	Scope	Condition
Detection rules	90+ rules across IT, OT and network-integrity threats, each mapped to MITRE ATT&CK or ATT&CK for ICS, or declared as a policy or anomaly check. A coverage view separates techniques WireTrace can detect from those observed, and exports an ATT&CK Navigator layer.	No coverage percentage is claimed.
Behavioural analysis	First-seen and rare destinations, new services, new internet communication, peer-group differences, scanning, lateral movement, beaconing, algorithmically generated domains, DNS tunnelling and unusual outbound volume, each explained against its baseline.	Deterministic analysis, not user behaviour analytics.
OT command monitoring	Unauthorised writes, parameter changes, controller mode changes and stops, program download and upload, forced I/O, firmware transfers, new or rogue masters and illegal function codes.	On the protocols listed in section 05.
Threat indicators	IPs, ranges, domains, URLs, file hashes and TLS fingerprints from public feeds and indicators you supply (CSV, STIX 2.1, or a commercial feed using your own key).	Air-gapped sites refresh through the offline data path.
Ransomware	Reconnaissance, lateral movement, mass file modification, ransomware extensions and ransom notes on file shares, and unusual exfiltration, correlated into one corroborated verdict.	Unencrypted SMB only; encrypted SMB sessions expose no file operations.
Files and web	File operations on network shares; transferred files hashed, matched to threat intelligence and scanned with YARA rules you manage; SQL injection and cross-site scripting attempts.	Unencrypted SMB and HTTP only. YARA coverage depends on your rules.
Network integrity	Adversary-in-the-middle techniques: TLS interception, TCP reset injection, ARP, DHCP and LLMNR poisoning, and DNS manipulation.	

Area	Scope	Condition
Vulnerabilities	Each device's identified OS, software, firmware and OT product matched to the NVD and OT vendor advisories, prioritised with CISA KEV and FIRST EPSS. Every finding states how it was matched, how confident the match is, and whether the vulnerable version is confirmed on that device.	Version-precise matches need a stated or collected version. IT patch verdicts need optional online sources.
End of life	Operating systems and products past vendor end of life, using lifecycle data that also works offline.	OT firmware lifecycle coverage is partial.
Exposure and risk	Risks & Exposures view (clear-text credentials, weak cryptography, exposed services, OT exposure); 0-100 Asset Risk Score (likelihood × impact) with a confidence indicator and factor breakdown; Risk Register with time-limited acceptance.	Asset risk, not enterprise risk management.
Investigation	One Alert Queue; investigations that gather evidence, build a timeline and show supporting and contradicting evidence; guided no-code threat hunting; on-demand packet capture from sensors.	Packet capture is on demand, not continuous retention.
Compliance evidence	Network observations mapped to 415 controls across seven frameworks: NIST CSF 2.0, ISO/IEC 27001, HIPAA Security Rule, IEC 62443-3-3, NERC CIP, NCA OTCC and NCA ECC.	About a third of controls can be evidenced from the network; auditors close the rest in the platform.

Baselines, segmentation and policy

A **Communication Baseline** records a window of observed traffic; operators review the communication pairs found and approve, monitor or deny each one. A **Protocol Behaviour Baseline** learns which fields and values each industrial protocol normally carries, and once approved, a new field or value such as a previously unseen function code raises a deviation. Deviation checks run on a schedule you attach.

Segmentation intent is declared as traffic rules and as IEC 62443 zones, conduits and target security levels. WireTrace continuously detects traffic that breaks that intent, including IT-to-OT conduit crossings and large Purdue-level jumps. **Policies** define what happens when conditions become true, and each runs as monitor only, recommend, require approval or automatic, with four-eyes approvals that expire, an emergency stop, change freezes, maintenance windows and a maximum number of assets per action. WireTrace provides the intelligence, intent and validation; your existing firewalls, NAC and switches enforce it.

08 Data handling

What is stored

Device profiles and the identity evidence behind them, with source and timestamps; classifications; communication records between devices; decoded protocol fields and flow records that support alerts; alerts, investigations, baselines, policies and the activity log. Full packet captures are stored only when someone records one on demand.

Who can see it

Access is governed by roles, and on a multi-tenant platform each tenant has its own assets, users and permissions. WireTrace detects account names and credentials exposed by clear-text and legacy authentication so weak practices can be fixed; revealing a captured secret is restricted to authorised roles.

Where it is stored

All storage and processing stay on customer infrastructure: the platform server and its sensors. Rumi runs on the same server with no external AI service. External intelligence feeds are optional.

How long it is kept

Retention periods are planned with the deployment, using the storage profiles in the Deployment and Sizing Guide. Compliance evidence can be exported as a signed evidence bundle for audits.

09 Integrations summary

Purpose	Integration
Security operations	Alerts to any SIEM over syslog (CEF, LEEF, RFC 5424/3164; UDP, TCP or TLS); native Splunk HTTP Event Collector output; QRadar and Microsoft Sentinel through standard syslog formats; JSON webhooks, email and SNMP traps.
Data access	Read-only REST API (OAuth 2.0 client credentials, scoped tokens, rate limits, IP allowlists) on your appliance, off until enabled.
Enforcement hand-off	Proposed block rules as text for iptables, nftables, Cisco IOS and pf; IP, domain and URL blocklists to pull; identity, classification and risk for a NAC through the API.
Inbound context and sign-in	Read-only SNMP, LLDP and CDP; device syslog and SNMP traps; LDAP/Active Directory and OpenID Connect single sign-on; STIX 2.1 and CSV indicators; vulnerability intelligence sources.

The WireTrace Integrations Overview describes each integration by direction and purpose.

10 Platform security

- **Trusted sensors only:** sensors are enrolled with an activation token, and sources that are not enrolled are refused.
- **Sign-in and roles:** LDAP/Active Directory, OpenID Connect single sign-on or an email one-time passcode; role-based access control governs what each user can see and do.
- **Accountable change:** every decision and action, administrative or automated, is recorded in an append-only activity log. Changes WireTrace makes can be time-bound and undone with a preview.
- **Controlled data exits:** the REST API is read-only and off until enabled, with scoped tokens, rate limits and IP allowlists. Audit evidence leaves as a signed evidence bundle.

Every content change is tested, peer-reviewed and versioned before release, and the Platform Updates page shows each intelligence area with its version and freshness.

11 Specification

Item	Specification
Platform	No agents on monitored devices. Customer-owned server or virtual machine; one self-contained installation package; resources sized from the host automatically
Sensors	Built-in sensor plus distributed sensors; receive-only capture interface and separate management interface; enrolled with activation tokens
Capture	SPAN, TAP, ERSPAN; sizing per the Deployment and Sizing Guide

Item	Specification
Active enrichment	Optional, off until configured; from the platform management interface; defined scopes and read-only credentials; OT-safe profile
Internet access	Not required; runs fully air-gapped, including Rumi
Analysis content	250+ protocols decoded; 600+ classification rules; 150+ applications; 90+ detection rules mapped to ATT&CK; 415 compliance controls across seven frameworks
AI assistant	Rumi, on the WireTrace server, with no external AI service; assistive, with click-through evidence
Authentication	LDAP/Active Directory, OpenID Connect SSO, email one-time passcode
Access control and audit	Role-based access; append-only activity log; segregated tenants
Updates	Release packages, including offline packages for air-gapped sites; connected appliances refresh vulnerability and threat data automatically; air-gapped sites through the offline data path
Resilience	Backup and restore

Limits, stated plainly

- WireTrace sees what reaches its sensors. Discovery and detection follow sensor placement, and devices behind routers appear by address unless their traffic is seen or enrichment adds detail.
- It does not decrypt traffic. Ransomware, file and web-attack detection work on unencrypted SMB and HTTP only.
- Model, firmware and serial number are recorded where devices state them or where optional enrichment collects them. Installed software inventory requires credentialed collection.
- Vulnerability findings state their match strength; devices that reveal no version cannot receive version-precise findings from passive monitoring alone.
- OT command monitoring covers the protocols named in section 05; other industrial protocols contribute identity and discovery.
- WireTrace defines, validates and monitors segmentation intent. Your firewalls, NAC and switches enforce it.

Review the architecture against your own network

A technical session walks through sensor placement, capture options and integration points for your sites, then a passive evaluation on one mirrored port shows what WireTrace finds in your environment.

Book a demo at wiretrace.io/request-demo or email sales@wiretrace.io