Catching hidden threats
Detecting evasive attacks with Corelight ML

These are the latest new methods that are designed to evade endpoint defenses:

- Lateral movement
- C2
- Exfiltration
- Credential theft
- Admin exploitation

But these actions cannot hide from the network

Corelight ML (machine learning) reads the evidence, detects the threat, and provides the explainability detection engineers demand

corelight

Table of contents

corelight

The threats your endpoint can't see

Modern adversaries are built to evade. They abuse built-in Windows functionality, tunnel through encrypted traffic, move east-west where EDR and firewalls lack depth, and weaponize the same admin tools your own team uses every day. Signatures catch yesterday's malware; they go quiet against credential abuse over Kerberos, an admin share silently accessed at 3 a.m., or a beacon hiding in DNS. When the endpoint is blind and the rulebook is empty, the **network becomes your last reliable witness.** AI only makes these threats faster, bolder, and more evasive and dangerous.

The threats your endpoint can't see

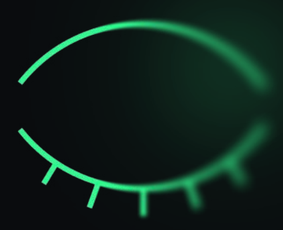

LOTL and credential abuse use native tooling endpoint sensors rarely flag

Signature gaps

Known-pattern matching misses novel C2, encrypted exfil, and evasive lateral movement

East-west darkness

Lateral movement and admin tool abuse between hosts happens where most tools never look

AI attacks

In the age of Mythos, AI finds hidden vulnerabilities, and provides new attack directions

corelight

Five types of evasive threats Corelight ML catches

Evasive attacks share a pattern: stay off the endpoint, blend into normal traffic, and keep moving. Corelight pairs two types of ML engines as part of a multi-layered detection architecture to break that pattern. Supervised models recognize invariant malicious behavior. Unsupervised anomaly detection learns what normal looks like in your environment, then flags the deviation. Together they help to cover the five categories that define a modern intrusion, and every detection they produce is one you can trust, tune, and investigate.

Lateral movement

The attacker is already inside and spreading

Includes behavior like anomalous RDP connection, SSH clients, connection pairs, admin file share access, executable file downloads

- Lateral Movement
- Execution

Command and control

The attacker phones home, quietly

Includes the use of HTTP C2 frameworks, infrastructure using multiple domains, NXDomain beaconing, DGA malware, connections to DGA domains, Tor, and malicious SSL certificates

- Command and Control

Data exfiltration

The data is leaving, disguised as normal traffic

Attacker use of exfiltration via DNS, a new anomalous service on an internet-facing server, a new anomalous service for internal communication, and using DNS reconnaissance preceding theft

- Exfiltration
- Discovery

Credential dumping & theft

The attacker steals the keys instead of breaking the lock

Includes east-west credential compromise via Kerberos use, of anomalous NTLM authentication, and anomalous admin share access tied to credential abuse

WHY IT MATTERS: Kerberos and NTLM abuse use built-in Windows functionality typically unseen by EDR

- Credential Access

Admin exploitation

The attacker uses your own tools against you

Attacker use of anomalous admin file share access from internal subnets, executable file transfers on a subnet, RDP connection, SSH client and connection pair

WHY IT MATTERS: Admin tool abuse looks identical to authorized activity; only behavioral baselining reveals what doesn't belong

- Privilege Escalation
- Defense Evasion
- Lateral Movement

corelight

How Corelight reads the evidence

Anomaly detection learns your network instead of guessing at it. Models train on unlabeled Zeek® data over a baseline period, then validates every anomaly against peer groups so a host's behavior is judged against its actual role. Detection runs on the sensor, so you add evasive-threat coverage without standing up cloud analytics or extra appliances. The payoff: fewer false positives, and alerts that arrive with the evidence already attached, ready to investigate.

1

Rich network telemetry

Zeek logs across DNS, SSL, SSH, RDP, SMB, NTLM, Kerberos, and connection data feed the models

- kerberos.log
- ntlm.log
- conn.log
- smb_files.log

2

On-sensor analysis

Supervised and unsupervised detection runs locally, no external connection required

3

Peer-group validation

An anomaly fires only when it is unusual to both the entity and its peer group. Anchors every alert to verified behavioral baselines and reduces false positives.

4

Evidence-rich alert

Detections surface with context logs showing behavior before and after the anomaly, ready for SIEM triage

31-day baseline training period before alerting begins

Nightly retraining so baselines track reality

Peer-group (collaborative-filtering) modeling to cut false positives

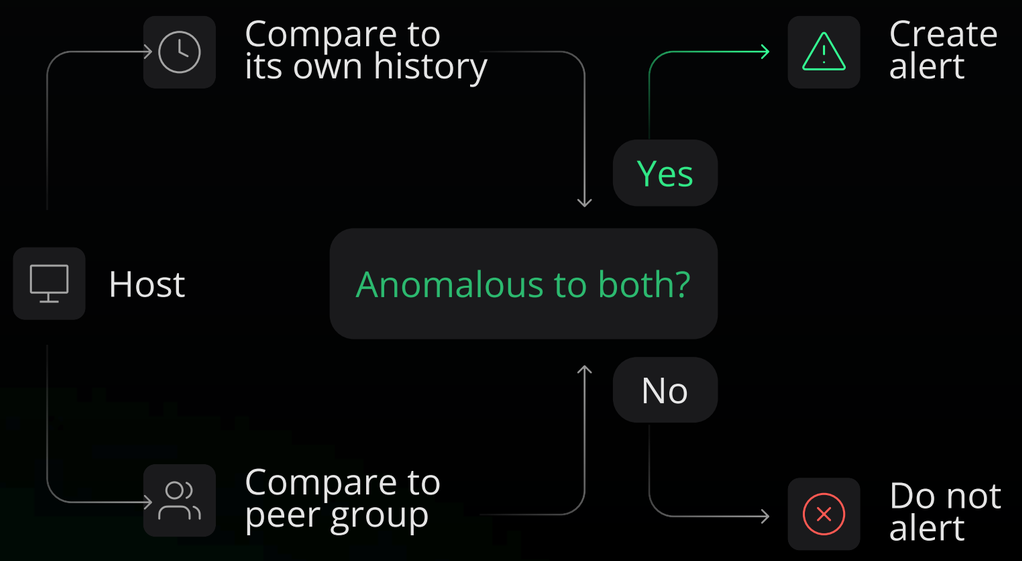

corelight

Trust, tune, and investigate

A detection you cannot explain is a detection you cannot defend. Corelight is built on open, inspectable Zeek logic, so you can see why a model fired, not just that it did. Peer-group validation filters activity that is unusual for one host but normal for its role, critical when distinguishing legitimate admin work from exploitation. And when something is noisy in your environment, you tune it. You do not just switch it off and hope. When something fires that matters, the evidence is already there.

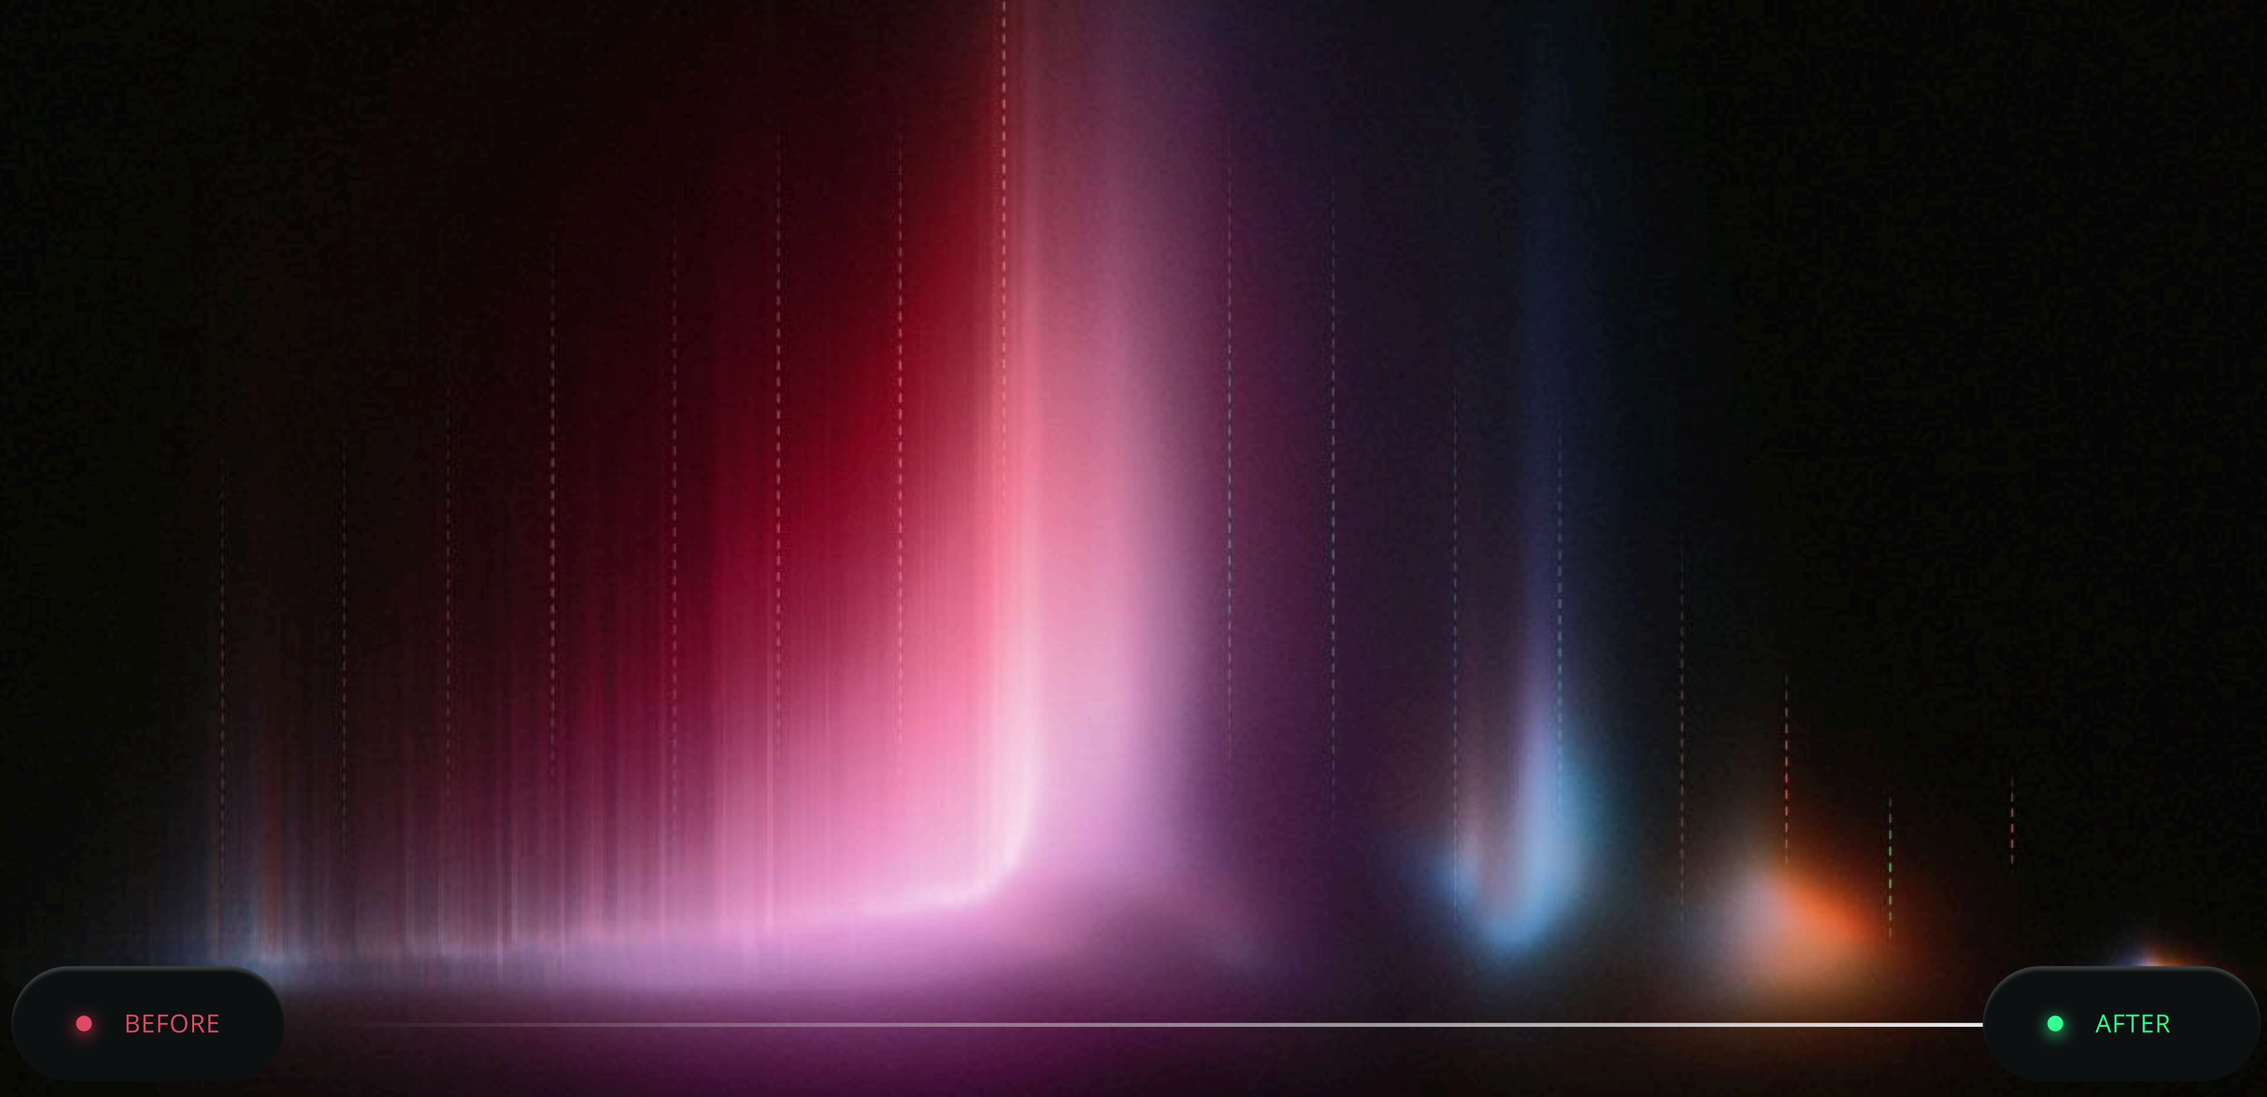

Trust it

Transparent Zeek-based detections with full logs and packet evidence – no black boxes, no guesswork

Tune it

Enable or disable individual models, set per-model thresholds, and adjust ignore lists to cut noise without losing coverage

Investigate it

Context logs show behavior before and after every anomaly, so triage starts with evidence, not questions

corelight

ML is one tool, not the whole strategy

Use the right tool for the job. Signatures and intel catch the known. Supervised ML and anomaly detection catch the unknown, the evasive, and the LOTL, including admin exploitation that looks identical to routine operations. Search lets you hunt what nothing else has flagged yet. Corelight bundles these layers on one platform with a shared evidence model, so your detections correlate instead of compete, and every alert, regardless of which layer fired it, is one you can trust, tune, and investigate.

- Signatures (Suricata, ET Pro)
- YARA file analysis
- Supervised ML
- Unsupervised anomaly detection
- Behavioral detection
- Threat intelligence (atomic IOCs)
- Threat hunting

Every layer correlates through a shared connection identifier, tying alerts back to the same network evidence for faster investigation.

corelight

The results

Less time researching, less time triaging, and AI-powered coverage of the techniques evasive attackers rely on, including the ones hiding behind your own admin tools.

97%

reduction in adversary and threat research time

97% reduction in adversary and threat research time—Launch Brief: Integrated Threat Intelligence & Open NDR Platform Detection Enhancements, p. 20

79%

reduction in threat triage effort

79% reduction in threat triage effort — Launch Brief: Integrated Threat Intelligence & Open NDR Platform Detection Enhancements, p. 20

3×

MITRE ATT&CK® technique coverage versus traditional IDS

3x MITRE ATT&CK technique coverage versus traditional IDS—Corelight FY26 OpenNDR Overview, pp. 46, 66

70+

additional MITRE sub-techniques beyond traditional IDS

70+ additional MITRE sub-techniques beyond traditional IDS—Corelight FY26 OpenNDR Overview, p. 46

corelight

The bottom line

AI is acting as a force multiplier for adversaries, accelerating and intensifying their ability to conduct lateral movement, C2, exfiltration, credential theft, and admin exploitation. The moves that hide from the endpoint cannot hide from the network. Corelight ML reads that evidence on-sensor, explains its reasoning, and hands you detections you can trust, tune, and investigate.

See how Corelight ML catches the threats EDR and signatures miss

Explore Corelight Open NDR

corelight