corelight

White paper

Accelerate insider threat detection with Corelight Open NDR

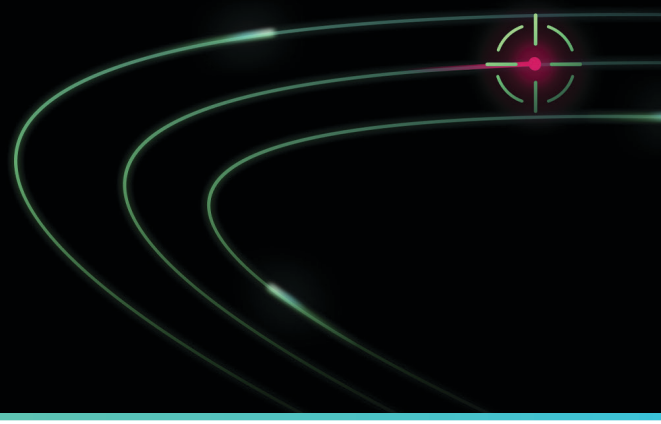

INTRODUCTION

Insider threats—compromised credentials, malicious employees, and negligent users—represent the attack surface that authorized access creates. Security controls based on perimeter defense and endpoints have a structural blind spot for threats already inside the environment. Insider incidents have the longest average dwell time of any threat category, and the damage is done before most organizations even know it has started. Existing controls cannot close this gap, not because they are poorly implemented, but because they were built for a different threat model.

CISA defines an insider as any individual with authorized access to an organization's network, systems, or data. The defining characteristic of an insider threat is that it exploits authorized access. The attacker has legitimate credentials, legitimate access rights, and the ability to take actions that are individually indistinguishable from normal work. An insider evades every control that identifies a threat based on known attacker behavior. The implication for threat detection follows directly from the nature of the problem. If the access is legitimate, the detection signal cannot come from the access itself. It must come from the behavior, and behavior is most reliably and most independently observed on the network.

INSIDER THREAT PROFILES: THE CISA TAXONOMY

CISA organizes insider threats into three categories. Understanding the operational profile of each is the basis for understanding how network evidence can accelerate insider threat detection.

Unintentional insiders

The largest category by volume, these are insiders who cause harm without malicious intent—misconfigured cloud storage, accidental data exposure, clicking a phishing link, and data handling policy violations without awareness of the rule. The harm is real regardless of intent; data is exfiltrated, systems are exposed, and the regulatory consequence does not distinguish between a malicious act and a careless one. NIS2, DORA, and SEC disclosure requirements apply to data exposure events regardless of how they occurred.

The detection profile is comprised of behavioral anomaly: unusual data access patterns, unexpected transfer destinations, and out-of-policy tool usage.

Malicious insiders

Individuals who deliberately exploit their authorized access. Three operational profiles of malicious insiders matter most for enterprise security teams:

- Data theft and pre-departure staging: An employee approaching resignation who begins staging and exfiltrating intellectual property, including code repositories, deal pipelines, and customer records, in the weeks before departure. The individual access events are legitimate; the pattern of behavior is not.
- Sabotage: A disgruntled insider—contractor, administrator, or recently terminated employee with residual access—who deletes, corrupts, or ransoms systems or data. Living-off-the-land (LOTL) tooling is the primary mechanism of sophisticated sabateur. Native operating system tools such as PowerShell, WMI, certutil,

and built-in scripting environments are used specifically because they blend into legitimate administrative activity that endpoint controls are calibrated to ignore.

- Espionage and nation-state recruitment: An insider operating on behalf of a foreign actor, over extended timeframes, exfiltrating high-value intelligence in low-volume increments specifically calibrated to avoid threshold-based detection. Dwell time in this category is measured in months or years.

Compromised insiders

An external attacker who has obtained valid credentials through phishing, credential stuffing, or social engineering and is now operating with the full access rights of the legitimate user.

Generative AI has materially lowered the barrier to creating compromised insiders. Highly personalized spear-phishing at scale, deepfake voice and video impersonation for remote onboarding fraud, and synthetic identity attacks are accelerating compromised insider attacks.

WHY EXISTING CONTROLS FALL SHORT

Effective insider threat detection requires monitoring agents deployed on every endpoint, every cloud resource, and every device with access to sensitive systems. Without that universal coverage, other security measures leave detection gaps that insiders can exploit.

EDR vendors recognize this. Their recommended approach to insider threat detection involves creating a behavioral baseline for each user, user group, job function, title, and device, then using a customized risk score per user and endpoint to surface unusual or suspicious activity. This is the right model in principle. The problem is that the model breaks down in practice on two fronts.

CORELIGHT OPEN NDR FOR INSIDER THREAT DETECTION

Corelight Open NDR focuses on what insiders cannot easily hide: their network behavior. Open NDR delivers comprehensive visibility into what users and devices are actually doing on the network, AI-powered multi-layer detection identifies risky behavior even when it seems legitimate, while agentic triage, combined with AI detection, speeds up insider threat identification.

Comprehensive visibility into insider activity

Insider threats almost always manifest in network behavior. Users access data or systems they do not normally touch. They move laterally across the environment. They stage and exfiltrate data. They use encrypted channels, VPNs, or LOTL tools that appear legitimate from any other vantage point. A detection strategy without visibility into these behaviors cannot support detection.

- **High fidelity:** Open NDR generates 70+ structured log types with deep per-protocol field schemas—dns.log (query name, rdata, TTL, transaction ID), http.log (full URI, user-agent, referrer, method), ssl.log (SNI, certificate subject, JA4 fingerprint, cipher suite), kerberos.log (client, service, error code, ticket flags), smb_files.log (file path, access commands), x509.log (full certificate chain). Every log is linked by a Unique Connection ID (UID) that ties the entire session together (connection metadata, application-layer activity, file transfers, and SSL handshakes), enabling analysts to pivot from a high-level alert to a complete reconstruction of the attack in a single query.
- **Encrypted traffic visibility:** Encryption is not a blind spot. Open NDR's TLS and cryptographic metadata—JA3/ JA3S/ JA4 fingerprints, SSH session metadata, VPN traffic analysis—enable detection of misuse and exfiltration in encrypted channels without breaking encryption. An insider routing exfiltration through HTTPS or an encrypted VPN is visible in the metadata even when the payload is not.
- **Static file analysis:** File transfer and extraction metadata are captured across protocols in Open NDR. Static file analysis with YARA rules detects when insiders move, share, or exfiltrate files that match sensitive data, such as proprietary source code strings, database headers, or documents marked "Confidential," as well as data staged data transfers using large compressed files.
- **Coverage across IT, cloud, and OT/ICS:** Open NDR extends visibility into operational technology and environments where an insider might attempt to disrupt industrial operations. Limited EDR deployment in these environments underscores the importance of specialized OT/ICS protocol coverage, as insider sabotage can have consequences beyond data loss.

Multi-layer detections mapped to insider TTPs

Insider threats rarely look like commodity malware. They blend into authorized activity. The detection strategy must leverage multiple detection engines. Parsing through dozens of network protocols, Corelight can identify the precise mechanisms of an attack as it unfolds.

- **Behavioral detection:** An insider trying to exfiltrate data can surface in network activity as tunneling, file movement, C2 communication, and other network

traffic. Motivated attackers can rotate domains to dodge an Indicator of Compromise (IOC) and modify payloads to evade a signature—but they cannot move laterally without generating SMB, RDP, or Kerberos sessions that Corelight Open NDR will parse, structure, and evaluate against learned behavioral baselines. Corelight operates over 100 behavioral detection models at the sensor level, analyzing how sessions behave rather than just what byte patterns they match. Examples of detected behavior include DCE/RPC lateral movement detection (identifying the specific Windows service invoked, from which account, to which endpoint), Kerberos attack detection (Golden Tickets, Silver Tickets, Kerberoasting— none of which produce a signature an attacker can study), DNS tunneling analysis (distinguishing covert C2 from legitimate high-volume DNS even when throttled), and SMB file activity classification (distinguishing reconnaissance from lateral movement staging).

- **Anomaly detection:** Defeating the negligent employee and the compromised account requires an intimate understanding of "normal." Open NDR integrates advanced anomaly detection as one of its multi-layered detection methods to establish a baseline of standard network activity for every user and device. When network traffic or a user's behavior deviates from established patterns—such as a user suddenly attempting to use a secure protocol to a secure server they've never accessed before—Open NDR flags the anomaly. This intelligent analysis is often correlated with an alert from another detection layer to help cut through the noise of daily operations, highlighting a true risk when a subtle behavioral shift occurs, indicating that a trusted account has been compromised or that an employee is acting recklessly.
- **Threat intelligence—IOC correlation against network evidence:** Even insiders communicate with the outside world. Corelight Threat Intelligence correlates high-fidelity indicators of compromise in real time and retrospectively. Detection scenarios include: an insider or compromised account beaconing to known malicious infrastructure; an internal host exfiltrating to domains associated with APT or criminal operators; a compromised session resolving C2 domains. Because IOC hits are validated against deep network context, false-positive rates are reduced, and insider behaviors are not buried in noise.
- **Storage exfiltration—producer-consumer ratio:** Open NDR's producer-consumer ratio calculations identify connections that indicate imbalanced, and possibly suspicious, data transfers. An insider staging bulk data to a personal endpoint would produce a directional and volumetric anomaly that neither firewall logs nor NetFlow can surface.
- **Cloud exfiltration:** Open NDR can detect a host exceeding the outbound bytes threshold and transferring data to AWS S3 storage within the time period defined by detection_window. An insider using personal or company cloud infrastructure to stage or exfiltrate data would be flagged so that appropriate containment actions can be taken.
- **DNS-based exfiltration and DGA detection**. DNS is a trusted, frequently under-monitored exfiltration channel. Detections that rely on VPC DNS logs are straightforward to evade. Monitoring raw north-south network traffic is the source of truth that is difficult to evade. Open NDR's dns.log captures every DNS query on the wire. Two distinct patterns are detectable: data exfiltration attempts appear as large, seemingly randomized DNS requests with data encoded in query strings; Domain Generation Algorithm (DGA) activity—used by compromised insiders running malware callbacks—is visible in the length and entropy of query strings.

Autonomous triage

Most vendors in this evaluation use AI to (merely) assist SOC analysts, including natural-language querying, AI-generated summaries, and chatbot-style investigation aids. These are useful features, but they still require the analyst to drive the investigation, interpret the results, and make the disposition decision. Corelight's Agentic Triage performs fully autonomous alert investigation. The agentic system identifies the riskiest entities, gathers and correlates network data associated with the entity, executes expert-designed security playbooks, investigates the alerts, renders a disposition (confirmed threat, benign, or escalate), and recommends specific mitigation steps. The analyst receives a completed investigation with forensic evidence laid out. Customers using Agentic Triage report that investigation times have dropped from 90 minutes to under 5 minutes.

For insider threat incidents, which are among the most time-consuming to investigate manually due to the need to correlate behavioral patterns across extended timeframes, multiple protocols, and identity context, the operational impact of agentic triage is significant in reducing triage times.

CONCLUSION

Insider threats are hard because attackers, whether unintentional or malicious, use legitimate credentials, have legitimate access rights, and can take actions that are indistinguishable from normal work. The advanced detection principle that closes the gap is behavioral analysis on immutable network evidence. Network telemetry cannot be disabled by an insider who lacks control over network infrastructure. It is not stored on systems that the insider can modify. It captures the behavioral deviations that authorized access cannot conceal—the access patterns that don't fit, the lateral movement that has no business justification, the exfiltration that breaks the expected direction and volume between two hosts.

Corelight Open NDR surfaces all three CISA insider threat categories, with comprehensive visibility, multi-layer detection coverage mapped to insider TTPs, and agentic triage capabilities that accelerate triage by 10x. We equip your SOC team with the authoritative truth needed to identify anomalous behavior, stop exfiltration in its tracks, and ensure that only the people building your business leverage your network.

To learn more about insider threat detection, request a demo at

https://corelight.com/contact

info@corelight.com | 888-547-9497