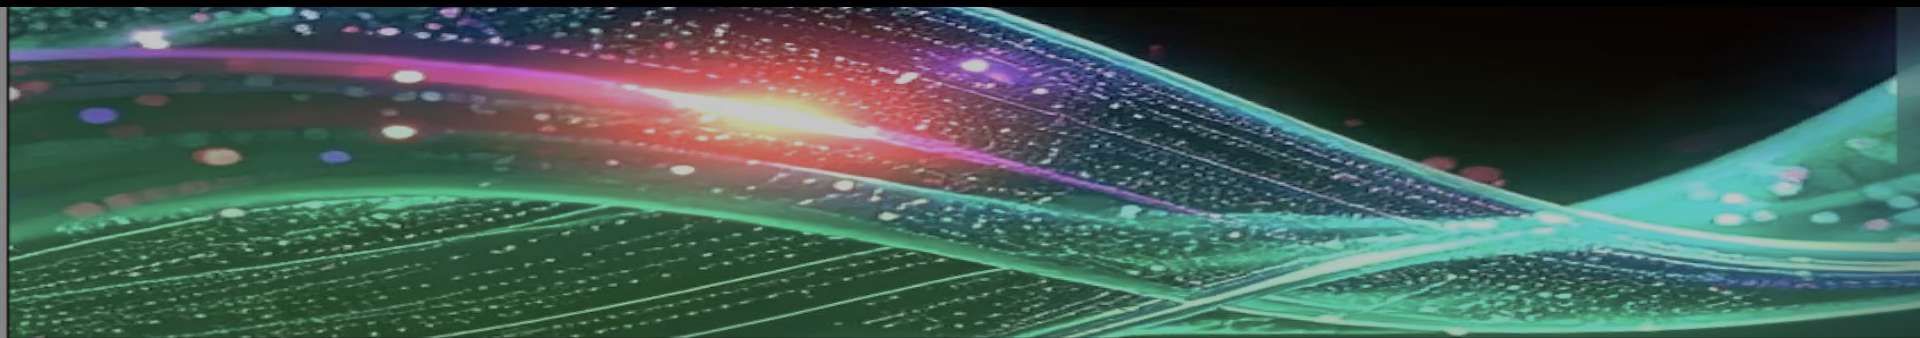




# Staying Ahead in the AI Cyber Race

*Illuminate and disrupt evasive threats with Corelight  
OpenNDR*



# Before we dive into the AI

The core of any defense is "The importance of network telemetry"

```
{ "_path": "conn", "_system_name": "sensor.violetbears.cyberlab.trycorelight.com", "_write_ts": "2026-03-17T16:16:03.572669Z", "ts": "2026-03-17T16:15:03.572516Z", "uid": "CfggWu3EsSH7ib9Ibk", "id.orig_h": "0.0.0.0", "id.orig_p": 270, "id.resp_h": "192.0.2.0", "id.resp_p": 514, "id.orig_l2_addr": "02:00:00:00:00:00", "id.resp_l2_addr": "fe:ff:ff:ff:ff:ff", "id.orig_l2_vendor": "unknown - Locally Administered MAC", "id.resp_l2_vendor": "unknown - Locally Administered MAC", "proto": "udp", "service": "syslog", "conn_state": "S0", "local_orig": false, "local_resp": false, "missed_bytes": 0, "history": "D", "orig_pkts": 1, "orig_ip_bytes": 214, "resp_pkts": 0, "resp_ip_bytes": 0, "tunnel_parents": ["CGPJzS337Pb1TTFhBi"], "spcap.url": "https://smartpcap.trycorelight.com/spcap/v1/?uid=CfggWu3EsSH7ib9Ibk", "spcap.rule": 1, "spcap.trigger": "all-unencrypted", "id.orig_h_name.src": "DNS_A", "id.orig_h_name.vals": ["qeqsag.com", "galyqaz.com", "23262ed0-598a-4a7b-bfff-5078a462f95e.local", "cb251212-263f-45b2-a013-d3147589e404.local", "puvytuq.com", "puzylyp.com", "pumyxiv.com", "qegyhigh.com", "5cb17307-a192-4e55-b27c-eafb9dbc9754.local", "7695117d-f76c-4511-abd1-8c62e9e2059e.local", "be592999-1476-4b98-af6b-8b60753aad6c.local", "puzywel.com", "6ebea2e2-52e0-4ca7-a9a7-72481f2c800f.local", "pufygug.com", "pumypog.com"], "orig_l2_addr": "02:00:00:00:00:00", "resp_l2_addr": "fe:ff:ff:ff:ff:ff", "community_id": "1:3Rcwg/BBqidXCAQdpinNzCo82LA=", "nat_orig_id": "b335b2e0f6e10f00ff39f9384b37ef2d0ab54db0", "nat_resp_id": "2bc727f74b8eff3bc02bad3a08b5be002c84b260" }
```

# About Corelight

- Enterprise class network detection and response platform
  - Forensic-grade network data
  - Complete range of detection capabilities from signature based and static files with threat intel to anomalous behaviour detection and AI driven detections
  - AI-assisted workflows
  - Tool consolidation with integrated IDS, PCAP, NSM and YARA for increased SOC efficiency
- Custodians of open source Zeek
  - Most widely used network security monitor
  - 25+ years of evidence and insights
- Gartner Magic Quadrant Leader (NDR)

Discover the technology quietly protecting the world's most elite networks:

**6/6**

US Military  
Branches



**\$10T+**

Financial  
Assets



**20+**

Critical  
Infrastructures



**16M**

Patient  
Records



# Technology of choice



Elite defenders are investing in Corelight and leveraging our technology to secure the products and services they deliver.

# Trends driving increased need for NDR in SOCs

## AI-driven threats and the operational complexities of defensive AI



Evasive, AI-driven  
Zero-Day Threats

SOC Alert Fatigue &  
Analyst Burnout

Regulations on  
sprawling  
infrastructure

Future ready AI SOC

# Open NDR platform AI capabilities

## AI-driven threat detection

AI/ML detections that are backed by forensic-grade network evidence - plus context - gathered in real time - on-premise, and in hybrid and multi-cloud environments.



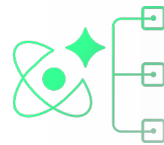
## AI-powered workflows

Expert-authored workflows to accelerate SOC workflows for triage and investigations with on-demand summaries, explanations, and response guidance.



## AI-enabled ecosystem

Providing structured, context-rich network data grounded in open-source standards already understood by LLMs, and designed to feed seamlessly into SIEMs and AI / ML pipelines -



# AI WORKFLOW

List Table Newest Detection

<b>Silver POST uri   172.27.0.137</b> 10 148 alerts from January 7th, 7:31pm to March 18th, 3:37am <span>Open</span> <span>Unassigned</span>	<div><span>Actions</span> <span>View Detection</span></div> <div><b>Description</b></div> <div><span>Generated from AI</span></div> <p>This Suricata rule is designed to alert on TCP traffic directed towards HTTP servers that produces a specific output from the <code>id</code> command, indicating a potential compromise. The rule checks for established connections and looks for the presence of the string <code>"uid="</code> followed closely by a numerical user ID. It also requires the presence of <code>"gid="</code> within 5 bytes and <code>"groups="</code> within 8 bytes of the previous content, confirming the output is related to user and group information. The use of PCRE (Perl Compatible Regular Expressions) further refines the detection to ensure the IDs are numeric and not just any characters. The rule is categorized as a major severity alert under the "bad-unknown" classification, suggesting that the identified traffic could indicate malicious activity.</p> <div><b>Detection Summary</b></div> <p>Status <span>Open</span></p> <p>Assignee <span>Unassigned</span> <span>Assign to me</span></p> <p>Number of Alerts 60</p> <p>First Alert Time February 17th, 4:55pm</p> <p>Last Alert Time March 18th, 3:33am</p>
----------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

*Historically, a junior analyst might stare at this and not know where to start*
















# INVESTIGATOR

## Highest-risk Entities ?

Entities are sorted by severity, then detection count, from the last 7 days

### Introducing Agentic Investigations ×

Maximize SOC efficiency with automated daily triage that investigates the last seven days of activity for your environment's top 30 entities. By running specialized queries and applying transparent expert-authored playbooks, Agent Lux empowers you to take decisive action with defensible logic.

Score	AI Insights	Entity	Locality	Highest Severity Category	Detection Severity	MITRE Tactics	Triage Status	Notes
10	 Likely Malicious	172.27.0.137	Internal	DetectSliver::SLIVER_HTT...	17 total detections 		0/17 closed	0
10	 Likely Malicious	172.27.0.138	Internal	DetectSliver::SLIVER_HTT...	16 total detections 		0/16 closed	0
10	 Likely Malicious	10.110.100.5	Internal	CORELIGHT MALWARE Ca...	16 total detections 		7/16 closed	1
10	 Likely Malicious	172.27.0.139	Internal	DetectSliver::SLIVER_HTT...	16 total detections 		0/16 closed	0
10	 Likely Malicious	192.168.100.10	Internal	ET MALWARE [CISA AA21-...	5 total detections 		0/5 closed	1

5 ▾ 1-5 of 18 items Page 1 of 4 < >



# AI WORKFLOW - UNDERSTANDING DETECTIONS

## ☰ Detection Summary

Alert Category	Type	MITRE ATT&CK Techniques	Number of Alerts	First Alert Time	Last Alert Time
ET ATTACK_RESPONSE Output of id co...	Suricata	-	60	February 17th, 4:55pm	March 18th, 3:33am

### 🔗 Rule Description

This Suricata rule is designed to alert on TCP traffic directed towards HTTP servers that produces a specific output from the `id` command, indicating a potential compromise. The rule checks for established connections and looks for the presence of the string `"uid="` followed closely by a numerical user ID. It also requires the presence of `"gid="` within 5 bytes and `"groups="` within 8 bytes of the previous content, confirming the output is related to user and group information. The use of PCRE (Perl Compatible Regular Expressions) further refines the detection to ensure the IDs are numeric and not just any characters. The rule is categorized as a major severity alert under the "bad-unknown" classification, suggesting that the identified traffic could indicate malicious activity.

### Suricata Rule

```
alert tcp $HTTP_SERVERS any -> $EXTERNAL_NET any (msg:"ET ATTACK_RESPONSE Output of id command from HTTP server"; flow:established; content:"uid="; pcre:"/\^d+[\^r\n\s]+/R"; content:" gid="; within:5; pcre:"/\^d+[\^r\n\s]+/R"; content:" groups="; within:8; metadata:created_at 2014_09_26, signature_severity Major, tag Description_Generated_By_Proofpoint_Nexus, updated_at 2019_07_26; classtype:bad-unknown; sid:2019284; rev:3;)
```

### Ask GPT

#### 🔗 What might this alert mean?

The Suricata alert you're seeing indicates a potential security incident involving an HTTP server's response to a command that reveals user and group information. Specifically, the alert triggers when a TCP connection from an HTTP server to an external network shows output related to the Unix/Linux `id` command, which discloses user ID (UID) and group ID (GID) details, as well as group memberships. The alert's criteria include specific content strings, such as `"uid="`, `" gid="`, and `" groups="`, along with regular expression patterns to validate the command's output format. The metadata indicates that the rule is marked with a "Major" severity level, highlighting its importance in identifying potential unauthorized access or misconfigurations that could lead to information leakage. Given the nature of this alert, it is crucial to investigate the source of the HTTP response and determine if there is any malicious intent or exploitation occurring on your network.

### Request More

# AI WORKFLOW - STEPS TO FOLLOW

🔍 What are some possible next steps to investigate this alert?

To investigate the Suricata alert for "ET ATTACK\_RESPONSE Output of id command from HTTP server," a Cyber Security Analyst would use specific Zeek logs to gather relevant information. Here's a breakdown of the Zeek logs to focus on and how to use them effectively:

## 1. HTTP.log

- **Purpose:** This log captures all HTTP requests and responses.
- **Key Fields of Interest:**
  - `ts` : Timestamp of the HTTP transaction.
  - `uid` : Unique identifier for the connection.
  - `id.orig_h` : IP address of the client making the request.
  - `id.resp_h` : IP address of the server responding to the request.
  - `url` : The URL that was accessed.
  - `status_code` : HTTP status code indicating the response.
  - `body_len` : Length of the HTTP response body.
- **How to Use:**
  - Look for HTTP responses that contain the `uid=`, `gid=`, and `groups=` strings in the body. This might indicate the output of the `id` command, which is indicative of a potential attack or misconfiguration.
  - Filter by the response status code, focusing on any non-200 codes, as they may indicate an error response that is relevant to the investigation.

## 2. conn.log

- **Purpose:** This log records connection-level data, including TCP connections.
- **Key Fields of Interest:**
  - `ts` : Time of the connection event.

# Win the AI Cyber Arms Race with Corelight

## Multi-layered AI-driven threat detection

- Broaden threat detection with supervised ML, unsupervised anomaly detection, behavioural analytics, and signature-based methods to cover over 80 adversarial techniques within the MITRE ATT&CK® framework.
- Identify risks across your environment by uncovering unauthorised access points, expired certificates, unauthorised protocol activity, and malware at scale in hybrid and encrypted traffic.
- Create custom detections for specific threats, with support for industry-specific protocols like Multicast and MPLS, as well as encrypted protocols such as SSL and IPsec.

## Expert-authored workflows combining AI and network context

- Accelerate incident response through native integration with existing security tools, including EDRs, firewalls, and IT ticketing systems.
- Accelerate time to understanding and enable high-confidence triage with guided response actions.
- Quickly baseline and compare historical data flows to identify behavioural deviations during an investigation.

## Quick access to data and support for compliance

- Strengthen control integrity, including access and network segmentation for Zero Trust compliance.
- Automate alerts and reports for supervisory tasks and compliance requests.
- Maintain detailed, easily searchable Zeek logs to simplify audit trails.
- Rapidly estimate and verify incident scope while improving governance via cost-effective data storage and retrieval.
- Extend look-back windows for deep forensics.

## AI-ready data for SIEMs and AI / ML pipelines - out of the box

- Support correlation in SIEMs or SOARs that now leverage LLM plug-ins or copilots.
- Feed logs directly into security data lakes and enrich them using LLM agents, without needing format conversions or APIs.
- Generate accurate explanations or correlations over raw Corelight logs—simply by uploading the files and prompting naturally.
- Build proof-of-concept AI copilots that ingest Zeek logs and generate narrative timelines, threat summaries, and IOC maps with contextual linking.

# Manufacturer Finds Massive Exfil Linked to a Zero-day

## The Customer



A manufacturer needed to investigate a critical asset for signs of compromise tied to a recently disclosed vulnerability in a widely used database platform.

The security analyst needed historical network visibility to audit the database's behavior and definitively determine if the system's integrity had been breached.



### Challenge

Existing server logs were insufficient to cover the extended timeframe of the suspected compromise, leaving the security team without clarity on a potential zero-day attack that may have occurred months prior.



### Solution

Using Corelight's network evidence, the analyst pulled logs from the suspected timeframe to determine whether any data had left the database.



### Results

Corelight logs confirmed that a substantial volume of data was exfiltrated to the internet, likely taking advantage of the vulnerability before it became public.

**Network evidence proved that tens of gigabytes left the building.**



Thank you

