White Paper

How Corelight delivered Zeek at scale for the DoD

Introduction

The best security practitioners in the world from hundreds of global enterprises—including Microsoft, Amazon, Department of Homeland Security, Department of Energy and Facebook—are strong advocates of the open-source Zeek network monitoring platform. Zeek transforms raw network traffic into high-fidelity data streams for incident response, threat-hunting, real-time analysis, intrusion detection, and forensics. Since all attacks must, at some point, cross the network, and hosts are often compromised, network data is a foundational source of truth for information security. Corelight Sensors significantly outperform open-source Zeek and provide rich, actionable network data (Zeek logs) that incident responders and threat hunters need to protect and efficiently respond to network events.

For over 20 years, the founders of Corelight have been building and improving the open-source Zeek software. Several years ago, they founded a company to provide a purpose-built, highly scalable sensor appliance to meet the demands of the most stringent security operations environments. Corelight radically simplifies the deployment of Zeek by providing an optimized, high-performance,

Comparison	Corelight Sensor	Open-source Zeek
Flexible data export	Yes	No
Hardware accelerated NIC	Included	Separate purchase and integration
Analysis throughput	VM / 2 / 10+ 25+Gbps	3–4 Gbps
3rd party integrations	Yes	No
Optimized file extraction	Yes	No
Filtering to control export volume	Yes	No
Comprehensive API and python client	Yes	No
Performance charts	Yes	No
Geolocation	Yes	No
Encrypted drives	Yes	No
Integrated security	Yes	No
Zeek package manager & custom scripts	Yes	No
Support	Yes, creators of Zeek	Community mailing lists
Staff required for deployment/integration	Minimal	Zeek experts and linux specialists
Updates and maintenance	Automatic, +optional	Manual patching, tuning, and updating
Time for deployment	15 minutes	Weeks to months

enterprise-grade solution. The table above highlights some of the differentiators Corelight's purpose-built sensor provides over open-source Zeek.

It is important to recognize that Corelight is not simply open-source Zeek on a server. Instead, the creator and core technologists behind the project carefully engineered a solution to run in the mission-critical environments where they have spent most of their careers. They developed numerous optimizations for higher performance and added dozens of new capabilities that are required by large enterprises and mission agencies. An example of this optimization: when open-source Zeek performs file extraction, it writes each file to disk every time it crosses the sensor. In contrast, Corelight Sensors automatically perform file de-duplication to conserve system resources and file storage.

Experienced security personnel understand that in order to efficiently investigate and respond to security threats, they need visibility—preferably rich, real-time data collection that addresses logging as part of a timeline and not a single atomic event—and to simplify data export to SIEMs and other analysis platforms. Corelight delivers native integration with multiple SIEMs, including Splunk and Elastic, and also provides content filters so customers can maximize analysis efficiency and manage SIEM data-processing costs. Ultimately, Corelight delivers the most powerful network visibility solutions for information security professionals, helping them understand network traffic to detect, stop, and remediate cyber-attacks and intrusions. The screenshot on the next page demonstrates just one piece of Corelight's powerful visibility: insight into DNS traffic that's frequently manipulated by attackers.[1]

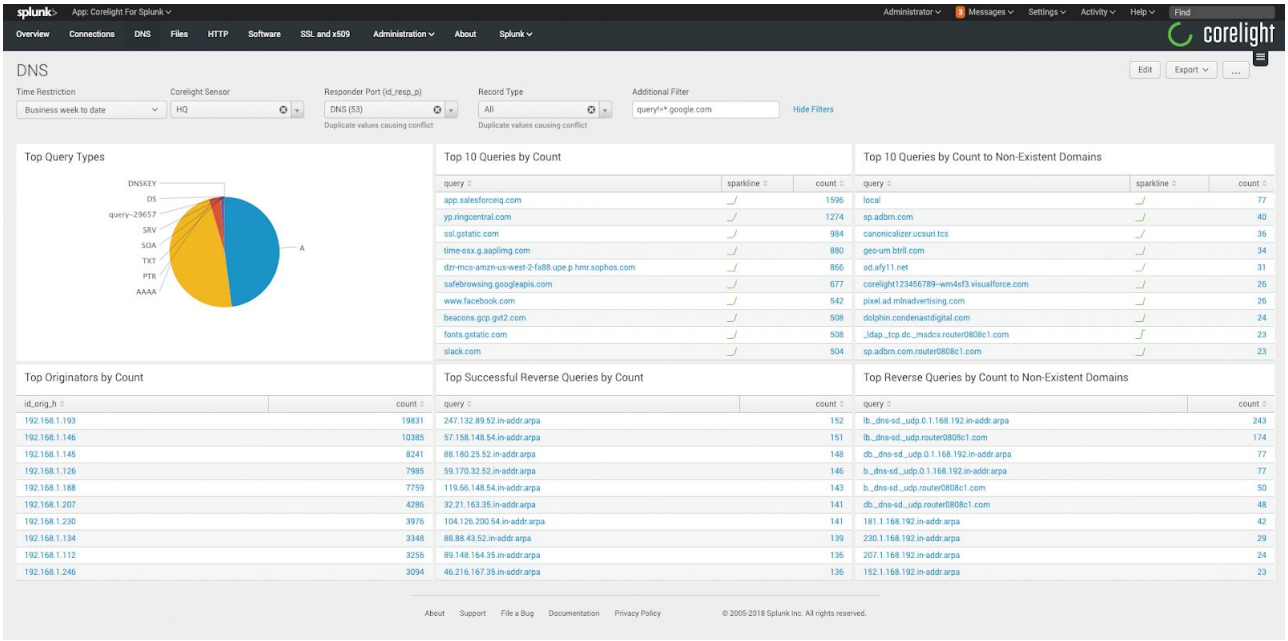

DoD pre-pilot results

The remainder of this document summarizes the results of a DoD proof-of-concept and pilot deployment of the Corelight AP 1000 Sensor that began in 2017.

The customer requested one sensor be shipped for lab performance testing using a traffic generator with four 1/10G SFP/SFP+ interfaces. The purpose of lab testing was to understand performance capabilities using a traffic profile based on current production traffic. This was a key first step with the customer as they already had an existing Zeek-based commercial solution deployed that could not scale and suffered packet-loss of up to 25%. The customer tested Corelight's hardware as well as commodity hardware running OS Zeek (provided by traditional networking manufacturers). The Corelight AP 1000 Sensor was the only solution that met the requirement of running Zeek at 10 Gbps, whereas the other commodity platforms dropped up to 25% of the traffic. Results demonstrated that the Corelight Sensor can process 30,000 CPS with a throughput of 16 Gbps and virtually no packet loss. Additionally, the sensor supported 20 Gbps bursts with less than 2% packet loss.

The table below contains additional performance results. As a result of the AP 1000's exceptional performance, a second Corelight Sensor was requested for additional performance testing to increase the throughput beyond a single sensor. During this second phase of testing, it was observed that 20% of the connections per second (CPS) were not valid security-related traffic. The customer implemented Corelight's advanced filtering to optimize performance by eliminating this traffic for analysis. After this performance improvement, the customer next wanted to determine how to best share Zeek scripts used in their operational environment (e.g., JA3)2 and the Corelight package manager addressed this request.

Traffic generator			Achieved								
CPS set	Conn set	Gbps set	CPS	Conn	Gbps	Corelight HTTP	SIEM HTTP	Corelight DNS	Ixia DNS	SIEM DNS	Packet Drops
20,000	2,000,000	39,200	18,000	2M	10	60,119	60,119	60,559	60,559	60,559	
30,000	2,000,000	39,200	27,000	2M	10	119,334	60,119	120,316	120,316	120,316	
30,000	2,000,000	39,200	27,000	2M	10	119,334	60,119	120,316	120,316	120,316	
30,000	2,000,000	39,200	27,000	2M	10	119,334	60,119	120,316	120,316	120,316	.036
30,000	2,000,000	39,200	27,000	2M	10	119,334	60,119	120,316	120,316	120,316	.005

During final pre-deployment design discussions, it was revealed that Corelight needed to implement further requirements to lock down the sensor. To fulfill this requirement, Corelight updated the GUI, implemented password complexity support,3 and added LDAP-S authentication.

The screenshots below depict the updated Corelight GUI. The first image is the Exporter section, which configures the streaming export of logs to platforms like Elastic, Kafka, and Splunk. The second screenshot shows the Statistics section, which provides insight into sensor health and performance metrics such as the volume of logs generated by the sensor.

Pilot

Following the successful proof-of-concept lab testing and Corelight's product updates, the sensor was successfully deployed. Key metrics that emerged from the Corelight pilot include:

- 60,000k exports per second (EPS) to Kafka
- A daily average of 1.8TB of Zeek logs ingested into Elasticsearch
- Packet loss of less than 0.05%
- Supporting sustained bursts of 10 Gbps

Soon after the deployment, Corelight released additional sensor models: the Corelight AP 200 (2 Gbps) and AP 3000 (25+Gbps). These new sensors enable customers to have right-sized visibility on smaller networks, remote sites, or large campus and data center networks.

Corelight also introduced optional flow shunting, a new capability for the AP 3000 Sensor which supports 40G QSFP+ monitoring interfaces. When enabled, the Corelight AP 3000 Sensor dynamically identifies specific long-running or bursty high-bandwidth connections and selectively bypasses such flows from analysis. Connection data and additional context are still monitored for shunted connections using the integrated FPGA NIC. In the case of well understood or uninteresting flows (e.g., TLS/SSL traffic), the analyst loses no context while greatly improving performance. With TLS traffic, for example, shunting occurs after the TLS handshake, so all clear-text data in the handshake is monitored prior to shunting. The encrypted packets and bytes are monitored, but the traffic is not analyzed. This technique provides context for TLS encrypted traffic sessions and greatly improves sensor processing performance.

As a result of the proof-of-concept, pilot test, and the analysis of alternatives (AoA) conducted by an FFRDC in July 2017, the customer determined that Corelight offers the only acceptable solution. Corelight successfully completed the pilot, which led to a sole-source Corelight acquisition.

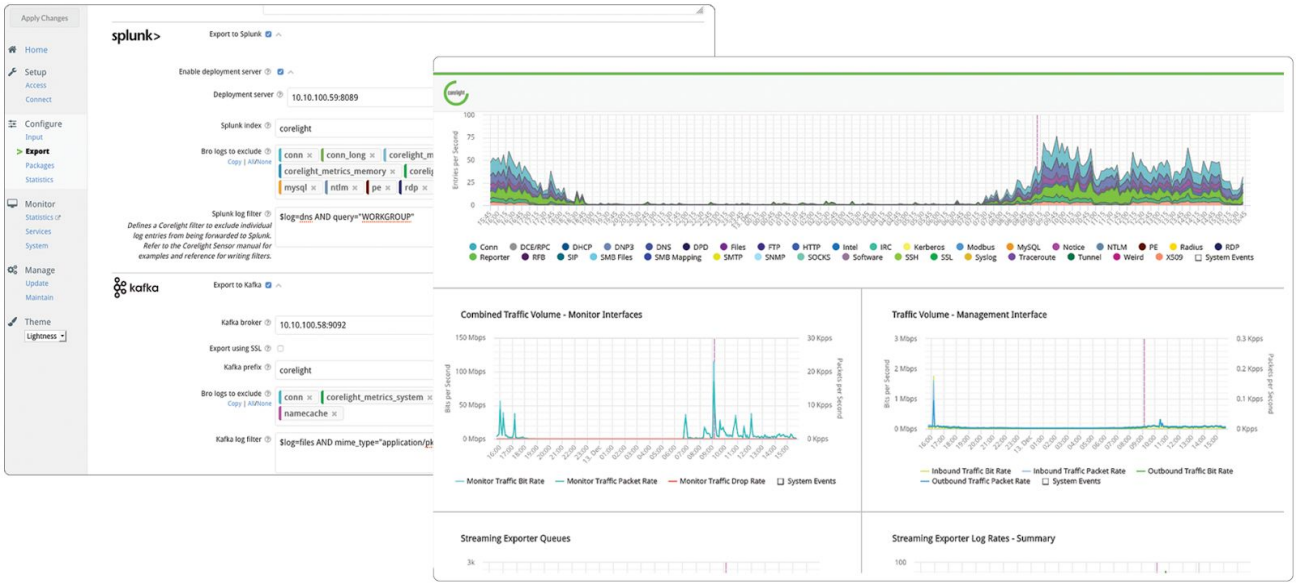

Corelight Sensor update
Corelight has since expanded its sensor family to include a cloud and virtual machine offering. The first implementation is with VMWare ESXi v6.5, and it supports throughputs up to 3 Gbps.

To round out the Corelight offering and simplify sensor management, Corelight also released a fleet manager. The Corelight Fleet Manager simplifies administration in distributed environments with multiple Corelight Sensors deployed. Fleet management builds upon the existing full-featured RESTful API to address three primary tasks:

- Initial setup and configuration
- Monitoring, remediating, and health
- Simplified multi-device management

Corelight Sensor software security certifications update
Corelight has completed the FIPS certification process and is prepared to begin the NIAP and DISA APL certification.

FIPS compliance
Acumen Security verified that the following product faithfully embeds a FIPS 140-2 validated cryptographic module, (26 June 2018):

- Corelight Sensor version 1.14.1

During the course of the review, Acumen Security confirmed that the following cryptographic modules are properly incorporated into the product:

- OpenSSL FIPS Object Module SE – FIPS 140-2 Cert # 2398.
- Red Hat Enterprise Linux libgcrypt Cryptographic Module v4.0 [1] and Red Hat Enterprise Linux libgcrypt Cryptographic Module v5.0 [2] – FIPS 140-2 Cert # 2657

Highlighted Corelight vs OS Zeek differentiators:

1. Process sustained network traffic exceeding 10 Gbps without packet loss
2. FIPS 140-2 compliant
3. Integrates with Kafka
4. Streams Zeek logs
5. Simultaneously transmits data to multiple analytic platforms (Elastic and Splunk)
6. Allows for one or more Zeek logs to be excluded from export
7. The AP 3000 Sensor supports custom shunting for large TCP flows
8. Corelight Sensors present Zeek log rates to estimate log volume
9. Implements a comprehensive RESTful API
10. Provides fleet management capabilities and a GUI sensor management tool

[1] The Corelight Splunk app is available at https //splunkbase splunk com/app/3884
[2] https //github com/salesforce/ja3/tree/master/bro
[3] The password requirement included a 15-character minimum length and required 1 upper / 1 lower / 1 number / 1 special characters

Defenders have always sought the high ground in order to see farther and turn back attacks. Corelight delivers a commanding view of your network so you can outsmart and outlast adversaries. We capture, interpret, and connect the data that means everything to defenders.

info@corelight.com | 888-547-9497