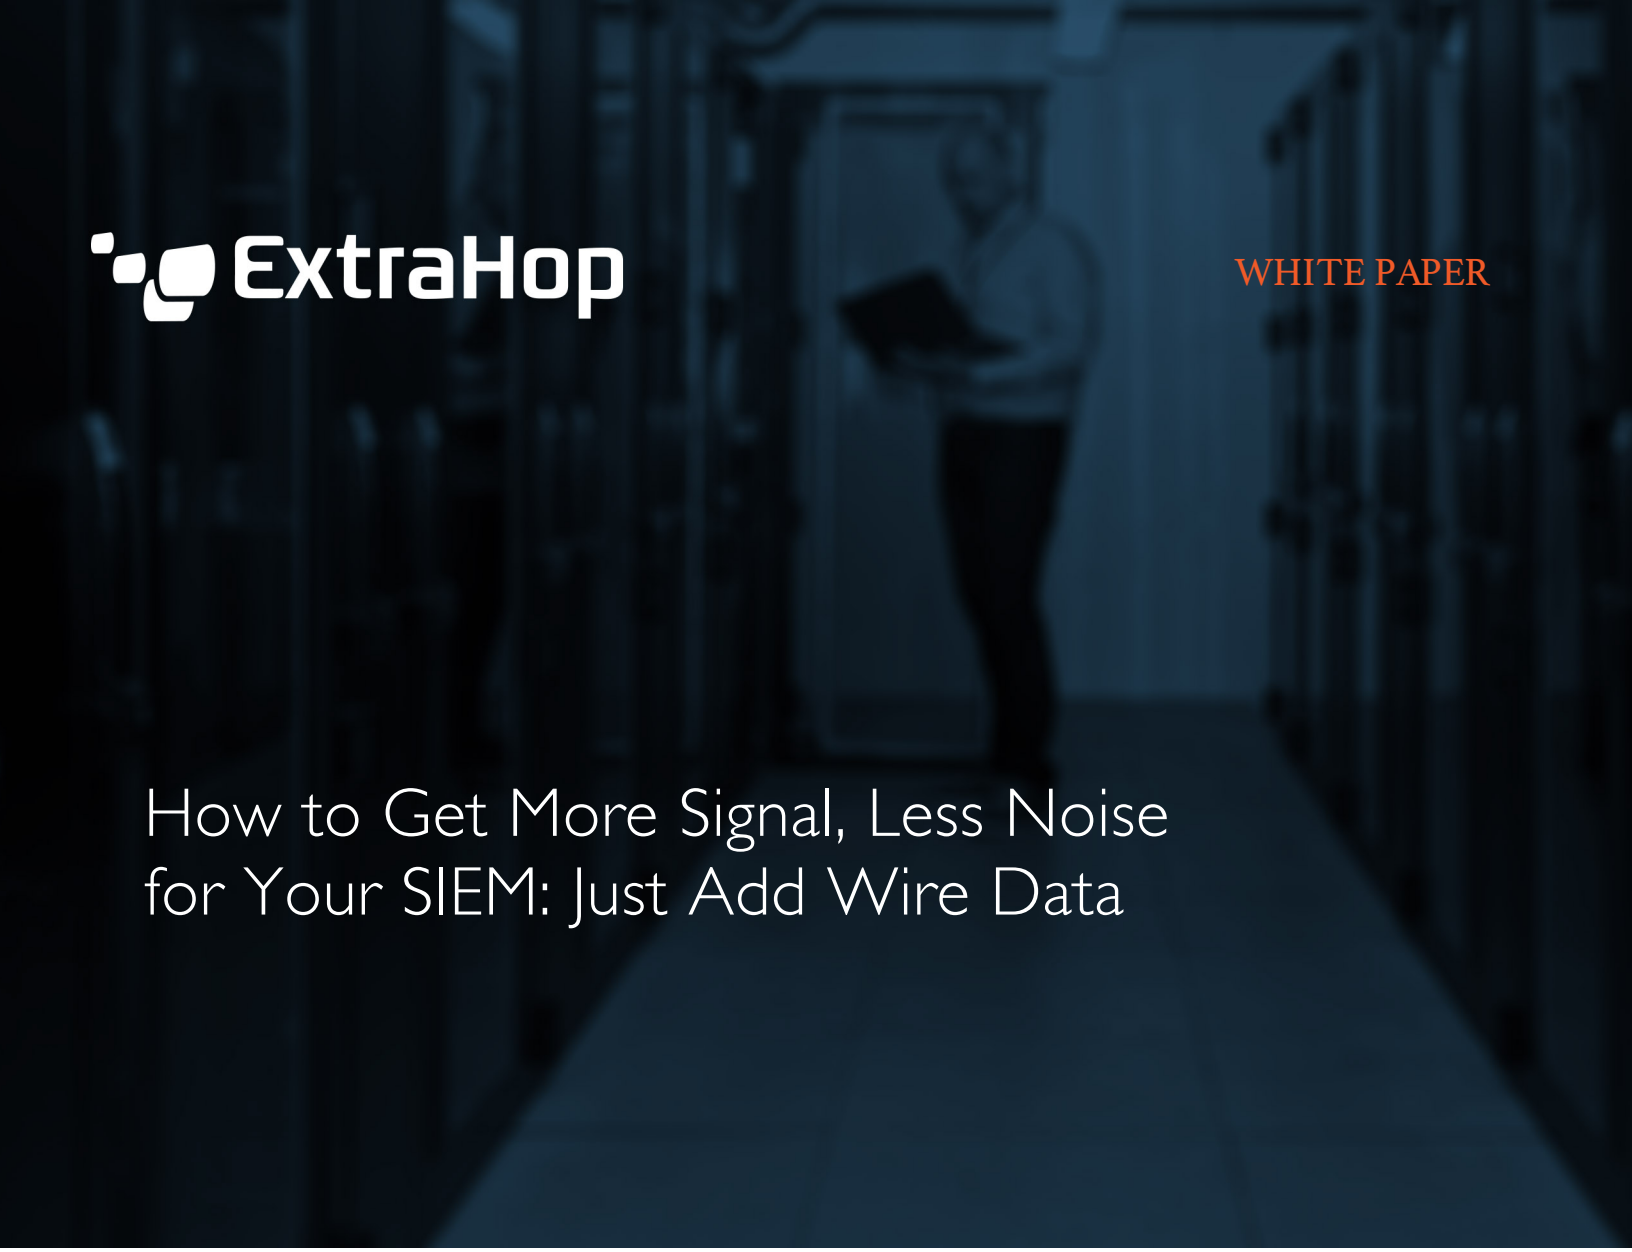

ExtraHop

WHITE PAPER

How to Get More Signal, Less Noise for Your SIEM: Just Add Wire Data

Abstract

Enterprise security organizations have made significant investments in SIEM platforms but struggle to optimize their accuracy. SIEM platforms depend on data from your environment to identify threats and enable investigation. Getting the right data is a tremendous task that is at the crux of every SIEM deployment's success. To improve the quality of their data sources, many enterprises have turned to the ExtraHop platform, which offers the ability to freely stream any network event or metric to third-party systems, including SIEM platforms.

This paper will explain how adding wire data from ExtraHop to your SIEM platforms will reduce the time spent by security teams in collecting and normalizing log data, minimize storage costs, and increase the effectiveness of your SIEM platform by providing better contextual information about threats and incidents.

The Future of SIEM Platforms Points to Wire Data

Over the last two decades, enterprise security organizations have amassed a broad portfolio of tools that generate an ever-increasing amount of data: Next-gen firewalls, IDS/IPS, end-point protection systems, threat intelligence platforms, and others. Organizations rely on security incident and event management (SIEM) platforms to make sense of this information and surface threats that would otherwise go unnoticed. SIEM platforms correlate logs, events, and alerts so that security teams can identify and investigate threats faster.

As critical as SIEM platforms are, most organizations would admit that they could work better. The ExtraHop platform introduces a rich new data source—wire data—to the SIEM platform that promises to simplify data collection, improve the signal-to-noise ratio, and increase their effectiveness.

What Is Wire Data?

Wire data is not just packets, but is the product of analyzing your data in flight, transforming raw packets into meaningful events and metrics that represent risk to your organization. The ExtraHop platform is the leading solution for real-time wire data analysis. Although primarily used by IT Operations teams, ExtraHop has also seen wide applicability for information security use cases, with dozens of real-world applications including data exfiltration detection, ransomware prevention, and continuous encryption auditing.

Examples of activity and behaviors included in wire data:

Application Behavior	Sensitive Protocols	Encryption	Compliance	Network Forensics	Vulnerabilities
Privileged user logins	Unencrypted FTP	Certificate expiration	SSH tunneling	Automatic discovery	Shellshock
Unauthorized outbound connections	Telnet	Key length	Non-standard ICMP	Precision PCAP	HTTP.sys
Lateral network traversal	Gopher	Outdated SSL sessions	Non-standard DNS	User activity	Turla malware
Brute force attacks	TACACS	MD5/SHA-1 cert signing	Non-standard HTTP	Network scanning	Heartbleed
Storage/DB access	SNMP v1, v2, v2c	SSL traffic by port	Disallowed file types	Triggers	FREAK SSL/TLS
Fraudulent transactions	Finger	Email encryption	Invalid file extension writes	Flow analysis	POODLE
	IRC	Wild card		Historical auditing	Logjam

SIEM Pain Points: It's All About the Data Sources

The complaints about SIEM platforms—that they generate too much noise and require ongoing management—can be addressed by improving the sources of data that you send to the platform. Wire data is cleaner, more comprehensive, and objective than log data and will dramatically increase the confidence you have in the quality of you SIEM results.

Wire Data Is a Better Data Source for Security Analytics

To date, SIEM platforms have relied on server and application logs, also known as machine data, for the contextual details needed for activities such as cyber hunting, ensuring compliance, and incident response. As helpful as this data is, IT and security professionals have learned to live with log data's shortcomings, as any security analyst who has had to collect and correlate data from multiple systems can attest.

Whereas logs are self-reported information, wire data is the observed activity on the network. This key difference makes wire data the ideal data source for SIEM platforms because it dramatically improves the signal-to-noise ratio, is always formatted consistently according to application protocols, and represents an empirical record of what actually transpired, not just what is reported.

	LOG DATA	WIRE DATA
Signal-to-noise ratio	Not all logs are created equal The details included in server and application logs depend on what the vendor or developer make available. Logs from different vendors may have different formats or include different levels of detail, which can prove challenging when trying to get a consistent, clean view of what is going on across heterogeneous environments.	Wire data has a higher signal-to-noise ratio than logs Real-time stream processing of network traffic produces structured metadata (formatted according to application protocols) about network activity the way that you want it: clean and consistent. There's less cruft that you have to sift through and pay for in terms of SIEM storage and indexing capacity.
Breadth of coverage	There are things you can log … and everything else Some important information is never included in logs, such as authentication success or failure messages contained in the HTTPS payload. In addition, you may not be able to install forwarders on some devices, such as IoT machines.	Every activity on the network produces wire data Real-time stream processing decodes all traffic on the network, regardless of the protocol or encryption. When you add programmability, you can extract all the information your SIEM platform needs at line rate without requiring logging.
Empirical observation	Attackers modify or erase logs Once an attacker has gained control of a system, they often erase or modify the logs to hide their activity. Essentially, this means that when you need log data the most, it can fail you.	Packets don't lie Security professionals have long understood that empirical evidence can be derived from the network. If a copy of all the network traffic is being passively analyzed by an out-of-band appliance, there's no way that attackers can hide.
Management overhead	Log data must be managed Security analysts spend a lot of time managing log data: deciding how to exclude duplicate events, ascertain the reason why logs are absent, and mapping data fields.	Wire data is automatically discovered and classified Real-time stream processing automatically discovers and classifies activity observed on the network, resulting in a clean and structured set of data.

Stream Any Wire Data to Your SIEM Platform

The ExtraHop platform can stream any wire data events or metrics to your SIEM platform, adding a rich new source of empirical data that complements and enriches your existing data sets. ExtraHop brings the richness of wire data to your SIEM platform, enabling you to fill in the missing gaps in visibility that you could not get with logs alone.

You choose what types of wire data you want to stream. For example, you can whitelist protocol communications for your applications and stream any policy violations—perhaps banned or sensitive protocols and services (FTP, telnet, gopher, etc.), database responses that exceed 10 MB, or how frequently specific login accounts are used and by which clients.

All of these activities and behaviors on the network can potentially represent risk to your organization. With ExtraHop, there is now a way to detect these behaviors in real time—*and stream those events and metrics to your SIEM platform* for alerting, correlation with other data, and historical auditing.

Open Data Stream

ExtraHop believes that you own your data, so we do not make you pay a "data tax" or lock you into a proprietary datastore. The ExtraHop Open Data Stream capability enables you to stream precise, real-time events and metrics from ExtraHop to your SIEM platform through syslog or a REST API. These events and metrics can be customized to match the SIEM platform. The Open Data Stream feature is free and enables you to make the most of the data that you already own through correlation with other data sets.

The ExtraHop community has documented integrations for Splunk and Sumo Logic, as well as other third-party security analytics systems such as the FireEye Threat Analytics Platform. The platform even supports the security data lake concept through MongoDB and Kafka feeds. Again, as you identify opportunities, you can replace or augment your log data with wire data to increase signal-to-noise ratio, minimize costs, and improve the effectiveness of your SIEM platform.

To learn more about the ExtraHop-Splunk integration, download the integration guide: *Integrate ExtraHop with Splunk*. This guide is representative of our integrations with other SIEM platforms.

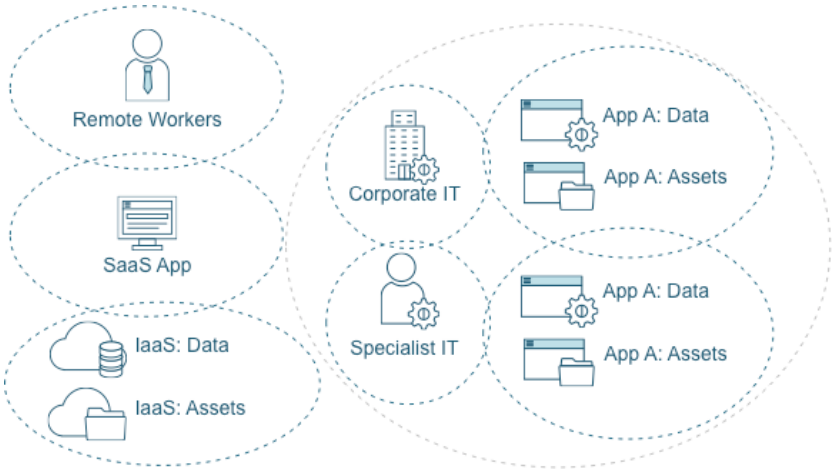

The ExtraHop makes it possible for InfoSec to "decentralize" security and equip each domain-specific team with the ability to define and monitor micro-perimeters for their own environment. You can also stream policy violations to your SIEM platform.

Benefits of Wire Data for SIEM Platforms

By streaming wire data into your SIEM platform and complementing or replacing log data where it makes sense, you will improve the effectiveness of your SIEM investment while simultaneously reduce management and infrastructure costs.

Speed Detection and Reduce Dwell Time

According to the 2016 M-Trends Report from Mandiant, attackers spend an average of 146 days inside network perimeters before they are detected.[1] During this time, they are stealthily performing reconnaissance, establishing persistence, and staging data for exfiltration. To identify bad actors faster, organizations need better visibility into east-west traffic, or lateral movement of data between servers.

No platform analyzes east-west traffic better than ExtraHop. While firewalls and IPS/IDS focus on north-south traffic entering or leaving the datacenter, the ExtraHop platform is positioned to analyze the growing amount of data travelling inside the datacenter. The Cisco Global Cloud Index estimates that east-west traffic within the datacenter will grow from 4.2 zetabytes in 2016 to 7.6 zetabytes in 2019, an annual compound growth rate of 24 percent.[2]

Traditional packet capture solutions cannot keep up with this challenge. At best, traditional packet capture will be used in forensic investigations, but it cannot proactively detect bad actors and alert InfoSec teams in time for them to take action. In contrast to legacy packet capture solutions, the ExtraHop platform is built for stream analytics, where analysis is performed on data in flight as opposed to data at rest. This enables ExtraHop to keep up with growing network traffic and provide visibility that enterprises need to detect malicious activity and behavior sooner.

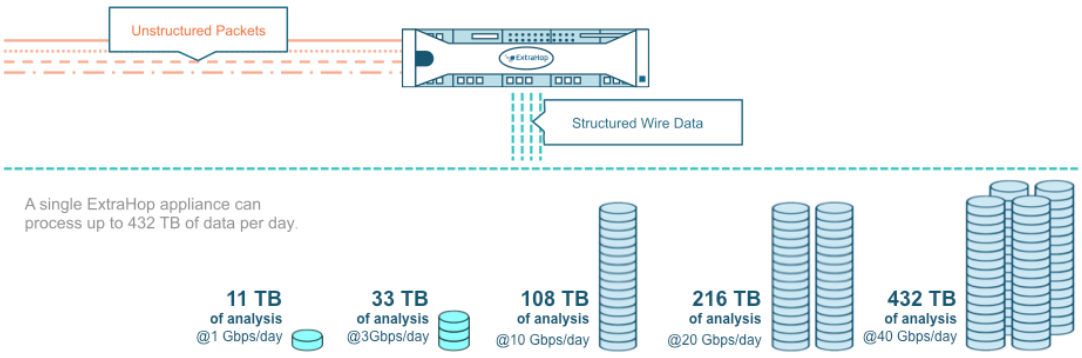

The ExtraHop platform performs stream analytics, transforming network packets into structured wire data in real time. This enables InfoSec teams to continuously analyze all east-west traffic and proactively detect bad actors.

Improve Investigation and Incident Response

One of the primary functions of a SIEM platform is to facilitate investigation of events and incident response. Analysts use the data stored in a SIEM to piece together the play-by-play of what happened (or didn't happen), including: the original attack vector and exploit used, the various systems compromised, the account logins used, the amount and type of data affected, and how much of that data left the environment.

Wire data provides the definitive record of what actually transpired leading up to and during an event or incident—and in a consistent and correlated format so that your analysts do not have to spend time putting together the digital forensic trail. Unlike logs, which are self-reported information that can be modified or deleted by attackers, wire data is the observed activity on the network that cannot be erased or altered.

[1] *M-Trends 2016 Report*, Mandiant: https://www.fireeye.com/current-threats/annual-threat-report/mtrends.html
[2] *Fifth Annual Global Cloud Index Report*, Cisco: http://www.cisco.com/c/en/us/solutions/service-provider/global-cloud-index-gci/index.html

The ExtraHop platform analyzes a copy of all traffic in real time to extract the events and metrics (authentication, database, storage, file transfer, and more) analysts need for investigation and incident response. ExtraHop also collects a complete and definitive record of the activity that is most needed during incident response, especially DNS, DHCP, and domain controller activity. When streamed into a SIEM platform, this empirical data enables your analysts to easily see all the communications between a particular IP address over a period of days, weeks, or months.

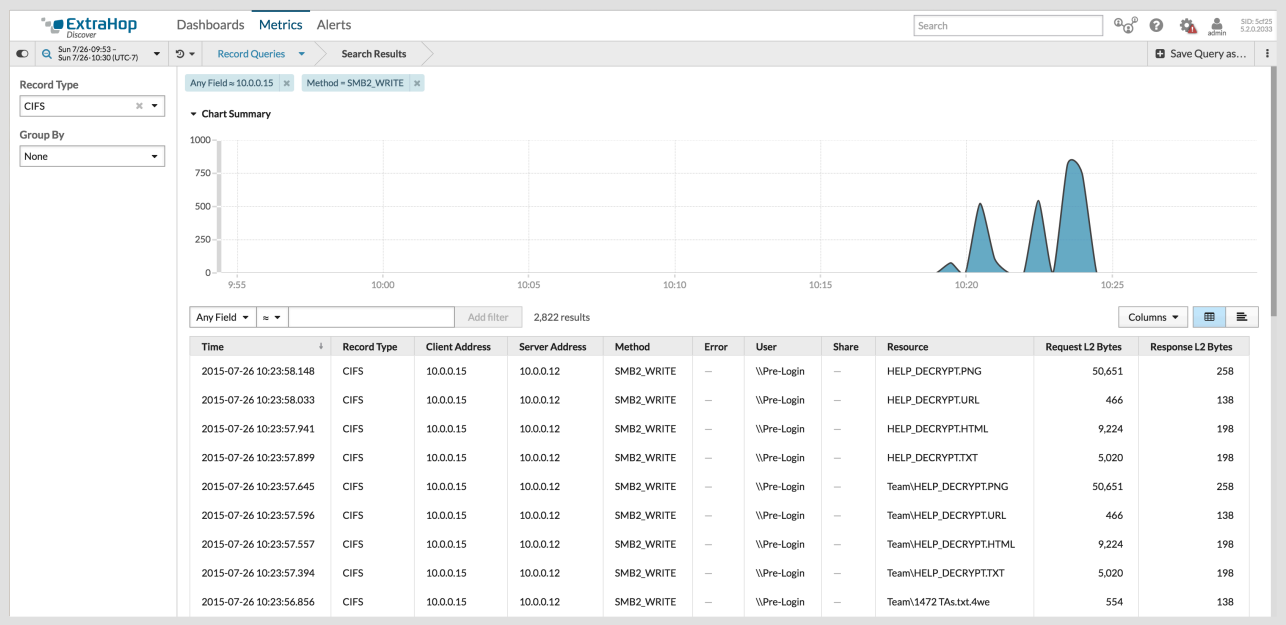

Time	Record Type	Client Address	Server Address	Method	Error	User	Share	Resource	Request L2 Bytes	Response L2 Bytes
2015-07-26 10:23:58.148	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	HELP_DECRYPT.PNG	50,651	258
2015-07-26 10:23:58.033	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	HELP_DECRYPT.URL	466	138
2015-07-26 10:23:57.941	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	HELP_DECRYPT.HTML	9,224	198
2015-07-26 10:23:57.899	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	HELP_DECRYPT.TXT	5,020	198
2015-07-26 10:23:57.645	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	Team\HELP_DECRYPT.PNG	50,651	258
2015-07-26 10:23:57.596	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	Team\HELP_DECRYPT.URL	466	138
2015-07-26 10:23:57.557	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	Team\HELP_DECRYPT.HTML	9,224	198
2015-07-26 10:23:57.394	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	Team\HELP_DECRYPT.TXT	5,020	198
2015-07-26 10:23:56.856	CIFS	10.0.0.15	10.0.0.12	SMB2_WRITE	–	\\Pre-Login	–	Team\1472 TAs.txt.4we	554	138

ExtraHop enables incident response teams to quickly search through a client machine's past transactions to uncover origination vectors and C&C communications. The screen above shows a filtered view of only CIFS WRITE transactions for a ransomware-infected client.

Simplify Historical Auditing and Trending

As mentioned earlier, wire data has a high signal-to-noise ratio, making it possible to store comprehensive records of activity for long periods of time. This observed record of activity assists with compliance reports and understanding historical trends. Of special interest is the authentication activity that ExtraHop tracks, including the use administrative accounts—63 percent of confirmed data breaches involved the use of weak, default, or stolen passwords, according to the Verizon 2016 Data Breach Investigation Report. [3]

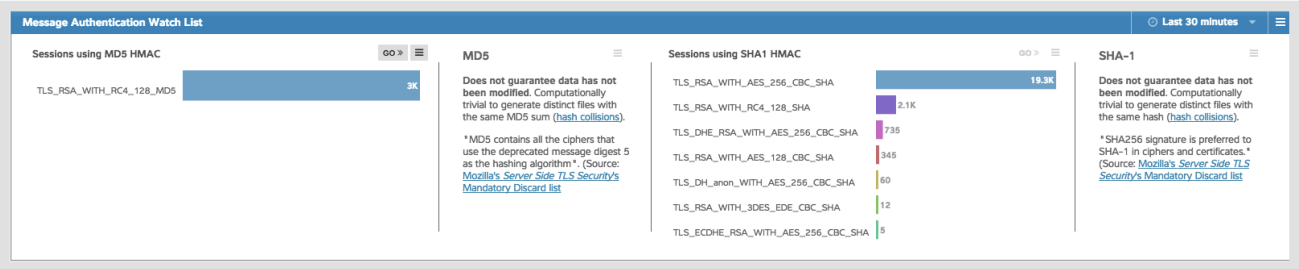

The ExtraHop platform makes encryption auditing simple by performing continuous SSL envelope analysis to determine the ciphers used, the expiration dates of keys, and other metrics required for compliance reporting. The screen above shows sessions using the non-secure MD5 and SHA-1 ciphers.

[3] *2016 Data Breach Investigation Report*, Verizon RISK: http://www.verizonenterprise.com/verizon-insights-lab/dbir/2016/

Reduce SIEM Indexing and Storage Costs

As you identify areas of visibility that are better suited for wire data, you can lower SIEM indexing and storage costs by cutting back on the amount of log data ingested into your SIEM platform. The ExtraHop platform also guarantees at least 30 days of lookback on each appliance with the option to store older metrics to your own network attached storage system. With wire data, there are no trade-offs between visibility and economics.

Additional Ways that Wire Data Enhances Security Infrastructure

Wire data is a critical missing piece to your existing security infrastructure. In addition to streaming events and metrics to your SIEM platform, the ExtraHop platform strengthens your current teams with:

- **Precise capture of digital evidence –** Set triggers to capture the digital evidence of policy violations, such as database responses larger than 10 MB or FTP connections to external servers. The ExtraHop platform offers the ability to capture the actual packets for precise flows based on custom-defined triggers. This facilitates the deepest level of event investigations as well as the digital evidence you need to prosecute or otherwise prove something happened.
- **Proactive management through integration with next-gen firewalls –** The ExtraHop platform offers generous APIs to enable proactive measures such as blocking IPs with a firewall or quarantining an infected client through network access control. Read about the ExtraHop-FireEye joint solution.

Conclusion

By adding wire data to your SIEM platform, you will simplify data collection, minimize infrastructure costs, and improve the overall effectiveness of your SIEM platform. Wire data offers several advantages over log data:

- A better signal-to-noise ratio with clean, consistent, structured data—saving you on SIEM indexing and storage costs
- Coverage for your entire environment using the network itself as a point of instrumentation
- Empirical observation of activity and behavior, unlike logs which can be modified or erased
- Elimination of data management challenges as stream processing automatically discovers and classifies all activity

The ExtraHop platform can provide the missing element needed to take your security infrastructure to the next level. With the ability to detect and stream real-time activity to other platforms that you already own, you can improve their effectiveness and ease of management.

ExtraHop Networks, Inc.
520 Pike Street, Suite 1700
Seattle, WA 98101 USA

www.ExtraHop.com
info@ExtraHop.com