



Threat Visibility for Cyber Hunters

By Sam Richman
Senior Security Systems Engineer

Abstract

Multiple branches of the United States military, spearheaded by United States Cyber Command, have embraced threat hunting as a way to defend against more sophisticated adversaries, as have an increasing number of commercial entities. This white paper explains how Cyber Protection Teams (CPTs) can use wire data to automate detection, speed investigations, and improve the granularity and collection of information. The paper includes examples of threat hunting workflows for rapidly investigating brute force attacks, data exfiltration, reconnaissance/lateral movement, ransomware infections, and malicious DNS behavior.

TABLE OF CONTENTS

Introduction.....	3
Using Wire Data to Hunt Threats	4
Automated Threat Detection	5
Active Hunting	7
Unified Traffic Visibility	7
Zero Knowledge Discovery of Assets and Traffic.....	7
Cross-Tier Protocol Visibility.....	7
Flexibility and Customization	7
Hunt Example 1: Brute Force Attack and Data Exfiltration Investigation	8
Hunt Example 2: Reconnaissance and Lateral Movement Investigation.....	10
Hunt Example 3: Ransomware Attack Investigation.....	12
Hunt Example 4: Russian DNS Queries and DNS Tunneling Detection.....	13
About the ExtraHop Platform	15
Conclusion	15

Introduction

The challenge of hunting bad actors, insider threats, and advanced persistent threats within an enterprise has increased exponentially as the IT landscape moves away from traditional datacenters and application architectures and towards hybrid and distributed environments comprised of highly virtualized and containerized assets. The sophistication of bad actors has also increased, reducing the value and timeliness of what self-reported data such as logs, SNMP, and NetFlow metrics can reveal about enterprise security.

The most effective method of detecting these sophisticated bad actors is a combination of automated threat detection and active hunting by Cyber Protection Teams (CPTs). Multiple branches of the United States military, spearheaded by United States Cyber Command, have embraced this strategy. Private industry has taken notice, and has dramatically increased investments in “hunt teams” in recent years. CPT operators are tasked with finding the proverbial needle in a haystack of petabytes of data generated by a multitude of heterogeneous assets communicating via numerous protocols.

What Is Threat Hunting?

Threat hunting starts with the assumption that bad actors have already breached perimeter defenses and are operating inside the environment. The goal is to proactively detect malicious activity by forming hypotheses about how attackers may have penetrated defenses, which systems are compromised, and what data they may have accessed.

Threat hunting efforts require familiarity with the environment, knowledge of potential weaknesses, and continuous collection of data. Therefore, only organizations with fairly mature security operations should formalize their threat hunting efforts. Organizations should prioritize securing their infrastructure and building out monitoring capabilities.

Security practitioners appreciate the idea of seeking out active threats instead of waiting until notified. In a 2017 survey of 330 cybersecurity professionals, Crowd Research Partners found that respondents spent much more time (43 percent of time) reactively investigating security incidents through activities such as alert triage than they spent proactively seeking out threats (only 22 percent of time). The same survey also found that 93 percent of respondents would rather work in a security operations center that focused on “lean-forward, proactive security capabilities.”



22%

Of time spent proactively
seeking threats

vs.

43%

Of time spent reacting to
security threats



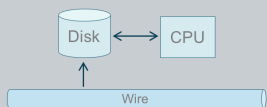
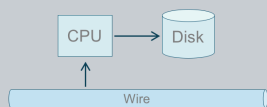
Early threat hunting efforts are paying off. In a separate 2017 survey of 306 respondents, the SANS Institute found that 91 percent of respondents improved the speed and accuracy of their response due to threat hunting, while 88 percent of respondents were able to reduce dwell time (the period from initial infection to detection), which is currently measured in months, a veritable eternity in computer time.^{1 2}

Using Wire Data to Hunt Threats

The industry is coming to the realization that network traffic is like the bloodstream of an organism³: just as with biological systems, symptoms of resident malware can be detected during the early stages of an infection. Real-time network traffic analytics is key to taking advantage of these early indicators of infection, since by the time vast quantities of stored network traffic are mined, these incipient infections may have already developed into a systemic and catastrophic disease.

Wire data serves as a trusted source of information because it is an observed record of activity rather than self-reported information such as log, SNMP, and agent data, and Therefore, wire data is highly resistant to compromise and can be used to validate security incidents or perform root cause analysis. In order to provide timely and actionable intelligence, however, this ocean of data must be mined in-flight and in real time. The traditional approach of writing petabytes of network traffic captures to disk and mining it after the fact is prohibitive in both time and cost. This approach simply does not scale in the modern reality of 40 Gbps and soon-to-be 100 Gbps networks.

ExtraHop takes the opposite approach, collecting raw network traffic and mining it in real time at 40 Gbps per appliance, automatically discovering client and server assets, and distilling terabytes of traffic per day into manageable and meaningful wire data. This is performed without the use of agents in a passive, out-of-band manner. Wire data is created by ExtraHop's real-time stream processor, which uses its native fluency in over 45 industry standard L2-L7 protocols to extract transactions and thousands of metrics from all these protocols simultaneously.

	Traditional Packet Capture	Stream Processing
How it works	Write to disk first, then analyze 	Analyze first, then write to disk 
Performance limits	Disk speed	Bus throughput and RAM
Lookback	Data typically stored for days	Data typically stored for months

This real-time approach allows CPTs to find anomalous activity in an efficient and timely manner, improving accuracy and response times between all teams involved in detecting and mitigating an attack. Real-time anomaly detection by the ExtraHop platform allows CPT operators to immediately shift focus to the methods and assets involved in an active attack without being overwhelmed by a

¹ SANS Institute, SANS 2017 Threat Hunting Survey <https://www.sans.org/reading-room/whitepapers/analyst/hunter-strikes-back-2017-threat-hunting-survey-37760>

² 2017 Mandiant M-Trends Report indicating a 99-day average dwell time by attackers.

³ Phys.org, Network traffic provides early indication of malware infection <https://phys.org/news/2017-05-network-traffic-early-indication-malware.html>

huge backlog of data. The insights contained within wire data would also otherwise require a time-intensive interrogation of multiple data sources such as firewall logs, Active Directory, web server logs, and more.

The comprehensive dataset created by the ExtraHop platform is available to CPT operators in an intuitive, visual user interface with a flexible workflow, allowing different teams or individuals to optimize the platform according to their needs. This intuitive user interface also has a low learning curve, allowing new operators to be effective in a short period of time with minimal training, especially valuable to CPTs with high turnover rates.

This paper will discuss how the ExtraHop platform fulfills two critical roles in threat hunting: automated threat detection, and active hunting by CPT operators.

Automated Threat Detection

While threat hunting is focused on human-driven activities, machine-driven analysis and data visualization can help to identify anomalous behavior that deserves special attention. With ExtraHop automatically discovering and monitoring network assets, the resulting real-time analysis of all transactions on the network (data-in-motion) allows CPTs to perform trending and alerting to detect, observe, and measure anomalous behavior seen on the network from any asset or user across all hosts, services, and transactions.

- Automatically discover, classify, and baseline all assets communicating on the network and discover their dependencies
- Identify and investigate anomalies by endpoint, protocol, or user
- Monitor use of banned ports, protocols, and services
- Detect reconnaissance and lateral movement behavior.
- Detect tunneled command and control and data exfiltration traffic
- Detect beaconing behavior
- Geolocation of traffic sources/destinations



Figure 1: Geolocation of traffic sources and destinations by protocol

ExtraHop utilizes machine learning to continuously build baselines for all clients, systems, applications, and infrastructure and then detect anomalies. These behavior-based alerts do not require any configuration by your teams. The ExtraHop platform builds baselines for new assets as soon as they are discovered by the system, providing continuous and complete coverage for dynamic environments.

Automatic anomaly detection provides your CPTs with a better understanding of what is abnormal in an environment, even if they may not have deep familiarity with specific applications. The alerts serve as starting-off points for investigation and include context to help staff determine the level of severity of the event, along with direct links into the asset detail page with the appropriate timeframe selected.

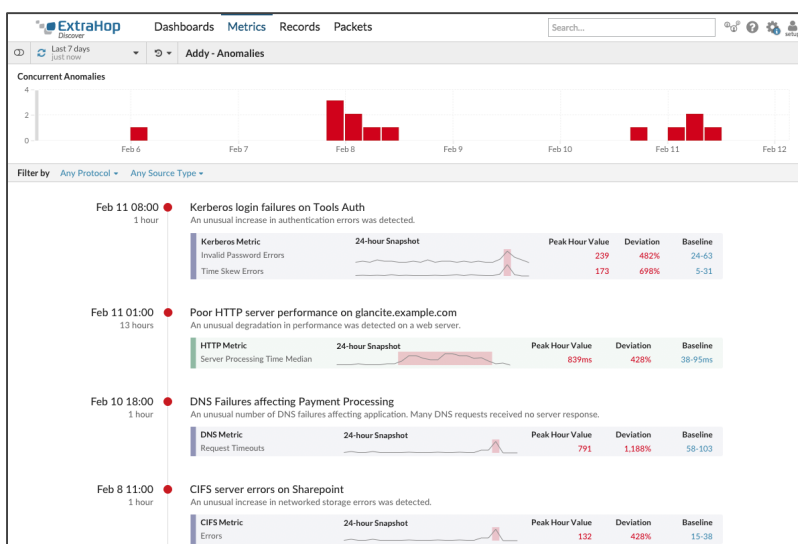


Figure 2: Anomaly detection detects abnormal behavior, such as a spike in authentication errors

Additionally, pre-built custom triggers written in JavaScript can immediately detect and take action on anomalous activity via alerts or REST calls to enforcement/mitigation platforms. These triggers can analyze any aspect of ExtraHop's supported network protocols, such as checking HTTP payloads for known malware packages, detecting indicators of ToR activity, and much more. As another example, the Scan Detection Bundle continuously analyzes network traffic for indicators of multiple reconnaissance methods and reveals them in real time, as shown below in Figure 3.

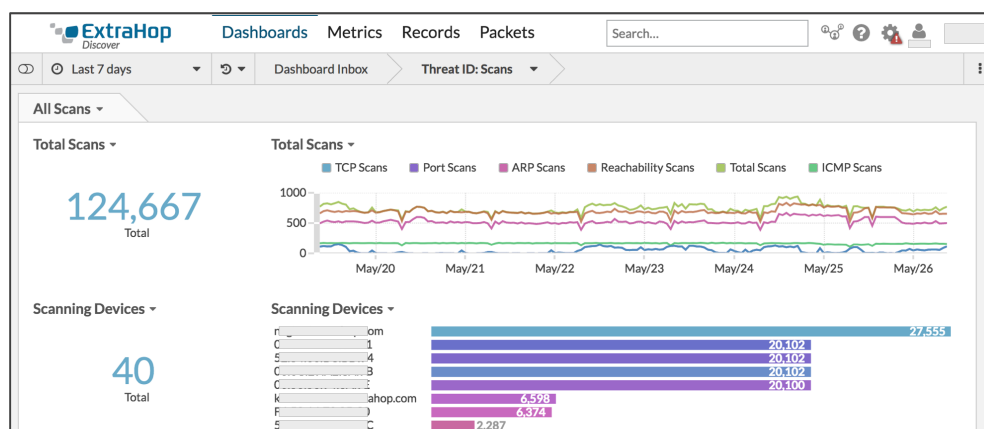


Figure 3: Scan detection dashboard

Active Hunting

ExtraHop is also used as an interactive detection platform by CPTs within networks that are suspected of being actively compromised or containing payloads associated with advanced persistent threats. ExtraHop may be installed as a permanent, resident tool within a network or can be included as small-form-factor or virtual appliances in the standard loadout for CPTs deployed to remote locations. ExtraHop's rapid, agentless deployment model makes it an ideal drop-in investigative platform.

The key characteristics of the ExtraHop platform which enable active hunting are described below, followed by four examples of active interrogation of the wire data that ExtraHop makes possible.

Unified Traffic Visibility

ExtraHop provides real-time visibility and analysis of wire data via its single platform workflow, from low-level packet captures all the way up to protocol transaction analysis and dashboard visualizations. Operators are able to easily pivot between PCAPs, metrics, and transactions within the same intuitive visual user interface, allowing them to achieve a complete incident timeline without the need to consolidate data from disparate tools.

Zero Knowledge Discovery of Assets and Traffic

Since ExtraHop does not require agents or foreknowledge of a network's architecture or activity, all assets transacting on the network, both known and unknown, are discovered. This posture readily reveals malicious/anomalous behavior without any modification to the environment under investigation other than providing access to a copy of the network traffic (SPAN, tap, or SPAN/tap aggregation solution). In zero trust environments, micro-segmentation policies can be readily audited from this strong security visibility posture.

Cross-Tier Protocol Visibility

CPT operators are able to easily pivot through all aspects of network traffic, moving from clients to servers, from one protocol to another, in order to follow the trail of an attack in an intuitive, visual workflow. Not having to learn a text-based query language has proven to be of great value, allowing even inexperienced analysts to quickly derive insight from the platform when on missions, as compared to other tools. Teams with high turnover also benefit from ExtraHop's low learning curve.

Flexibility and Customization

Easily created dashboards target specific protocols for monitoring, such as DNS, storage, SSL/TLS, HTTP, and more. Operators with varying experience levels can readily create dashboards on demand within the visual user interface, as well as make adjustments to existing dashboards with minimal effort and training. This flexibility, combined with separate accounts for each user, allows teams to operate effectively even without the assistance of senior operators, thus allowing more teams to be deployed concurrently and with a high degree of autonomy.

- ExtraHop's JavaScript-based trigger engine allows for a high level of custom detection as described earlier, such as identification of beaconing behavior, DNS tunneling indicators, anomalous user behavior, and more.
- All collected wire data may be extracted from the ExtraHop platform via its REST API, or streamed in real time via its Open Data Stream to industry standard big data platforms for centralized collection or correlation with other data.
- ExtraHop integrates with any platform offering a REST API, including security orchestration platforms like Phantom and ServiceNow.

Let's walk through four examples of threat hunting using the ExtraHop platform.

Hunt Example 1: Brute Force Attack and Data Exfiltration Investigation

A CPT operator is able to easily investigate detected anomalies or hunches by using ExtraHop's Live Activity Maps, a dynamic and real-time visualization of asset relationships and traffic patterns. In this example, a brute force attack against an internal database server was automatically detected, and indicated on the map in orange. (Note that another alert condition is also indicated on the map in red which can be investigated later):




Figure 4: Live Activity Map showing a brute force attack detected

One click focusses our attention on the database server under attack, and reveals all other traffic involving the suspect server, which includes outbound NTP traffic destined for a border router, a violation of security policy. NTP tunneling is a subtle and effective method of data exfiltration¹ because NTP traffic is often overlooked in an enterprise. However, using wire data analytics, just a few clicks have revealed a potentially serious security breach. A quick glance at an NTP security dashboard confirms that the NTP traffic destined for the unapproved external server is invalid, providing further evidence of tunneling behavior.

¹ Dark Reading, Simulated Attacks Uncover Real-World Problems in IT Security, <https://www.darkreading.com/cloud/study-simulated-attacks-uncover-real-world-problems-in-it-security/d/d-id/1330553>




Figure 5: Brute force attack investigation in progress

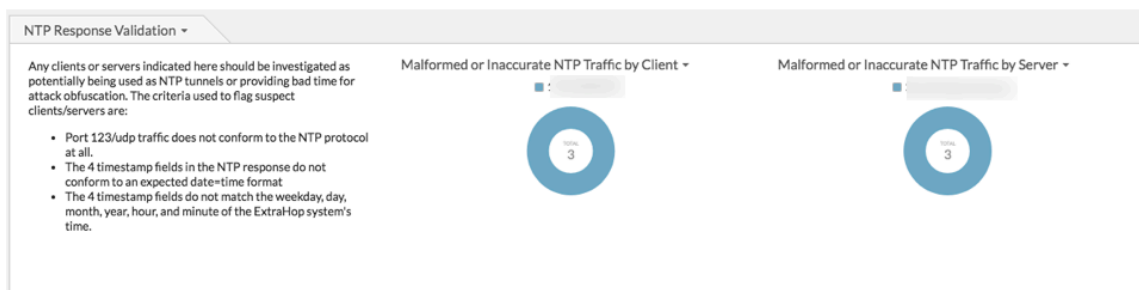


Figure 6: NTP Security Dashboard

Continuing the investigative workflow, clicking on the link between the DB client and server reveals the actual transactions occurring, indicating a series of failed root logins, followed by successful database commands. Because these database transactions are extracted passively from network traffic, they cannot be altered or erased by the bad actor.

Any Field ▾ ▾ ▾		Add filter		63 records				Fields ▾	
Time	Record Type	Client IPv4 Address	Server IPv4 Address	Method	Error	Database	Statement	User	
2018-01-02 15:22:50.1...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:51.8...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:53.1...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:54.1...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:55.2...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:56.2...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:57.2...	DB			OTHER	#28000Access denied for user 'root'@'...	--	--	root	
2018-01-02 15:22:59.7...	DB			OTHER	--	--	--	root	
2018-01-02 15:22:59.7...	DB			SELECT	--	--	/* mysql-connector-java-5.1.44 (Revision...	root	
2018-01-02 15:22:59.8...	DB			SELECT	--	--	SELECT version()	root	
2018-01-02 15:22:59.8...	DB			SELECT	--	--	SELECT DATABASE()	root	

Figure 7: Brute force attack database transaction records

We now have a deep understanding of the sequence of events in this attack from the initial brute force attempt, to a successful database infiltration, and finally to data exfiltration via NTP.

Hunt Example 2: Reconnaissance and Lateral Movement Investigation

A CPT operator is also able to react to automated detection of suspicious activity and directly pivot into an incident investigation/response workflow within the ExtraHop interface. In this example, Addy has detected evidence of reconnaissance/lateral movement outside of the normal baseline for this enterprise:

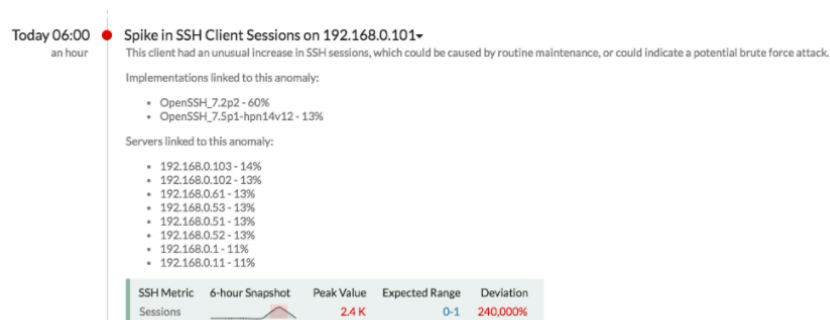


Figure 8: Addy alert for reconnaissance/lateral movement detection

Clicking into the anomaly reveals a Live Activity Map of the traffic patterns observed:



Figure 9: Reconnaissance/lateral movement live activity map

Note the *files-internal* server (also in red) which has been accessed via CIFS as part of this incident. The cross-protocol nature of wire data allows an investigation to start from SSH and seamlessly surface other suspicious protocol activity. Clicking on the connector between the file server and the client reveals the CIFS transactions, showing that sensitive files are being accessed:

Client = laptop489.extrahop.com		Server = files-internal.extrahop.com							
Any Field		Add filter		133 records		Fields			
Packets	Time	Record Type	Client IPv4 Address	Server	Server IPv4 Address	Method	Error	Resource	User
1	2015-07-26 21:48:19.1...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_WRITE	—	corporate\updated_expansion-opportunities-...	\\Pre-Login
2	2015-07-26 21:48:19.1...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_SET_INFO	—	corporate\updated_expansion-opportunities-...	\\Pre-Login
3	2015-07-26 21:48:19.1...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_SET_INFO	—	corporate\updated_expansion-opportunities-...	\\Pre-Login
4	2015-07-26 21:48:05.9...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_WRITE	—	—	\\Pre-Login
5	2015-07-26 21:48:05.9...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_SET_INFO	—	engineering\internal-only-branch-strategy-16...	\\Pre-Login
6	2015-07-26 21:48:05.9...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_SET_INFO	—	—	\\Pre-Login
7	2015-07-26 21:47:48.7...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_WRITE	—	—	\\Pre-Login
8	2015-07-26 21:47:48.7...	CIFS	192.168.0.101	files-internal	192.168.0.53	SMB2_SET_INFO	—	finance\updated_invoice-12603.gif	\\Pre-Login

Figure 10: Lateral movement CIFS transaction records

Clicking back to the Live Activity Map and adding one additional hop reveals what other assets have been interacting with the *files-internal* server, including CIFS clients which could be infected with malware deposited on the file server. Attacks are often multi-pronged, and can include both data exfiltration of sensitive files as well as implantation of malware.

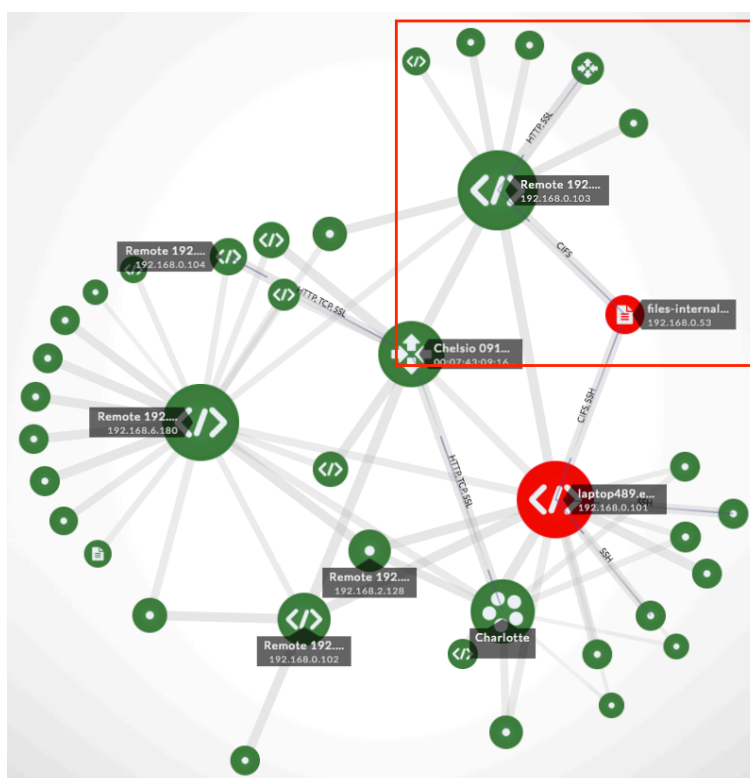


Figure 11: Expanded live activity map for investigation

Hunt Example 3: Ransomware Attack Investigation

The CIFS protocol-level visibility and cross-tier correlation discussed in the previous example is also foundational to ExtraHop's ability to detect ransomware attacks in real time, and to determine the source of the malicious payload which initiated the attack. In the following dashboard, clients performing file WRITE/MODIFY operations with suspicious file extensions are revealed.

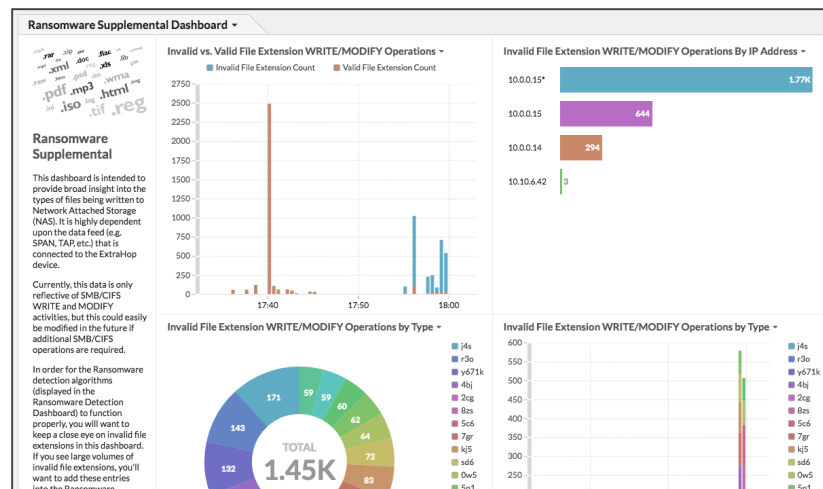


Figure 12: The ransomware detection dashboard reveals infected clients, including 10.0.0.15

Three clicks more result in a query of the CIFS transaction records for one of the client IP addresses identified as performing ransomware-like behavior (10.0.0.15) reveals the "HELP_DECRYPTING" files that the ransomware package created, shown in Figure 7 below.

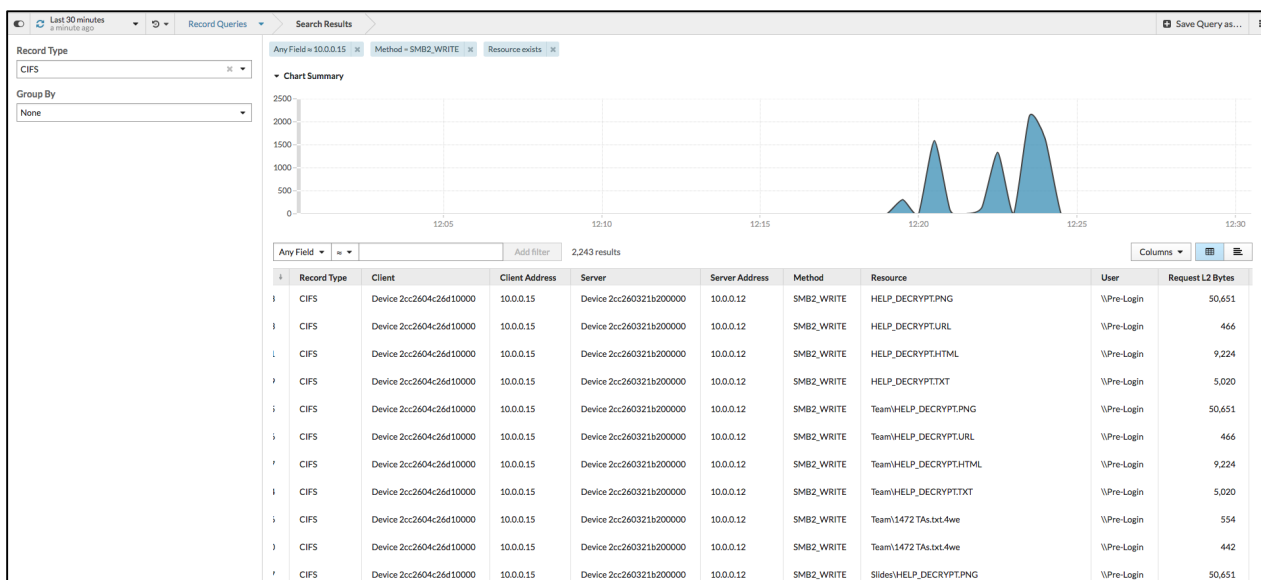


Figure 13: The ransomware package's component files on 10.0.0.15

One more search for the client IP address reveals all HTTP URIs accessed by this client in the same timeframe and provides an investigative path to determine how this client became infected with ransomware, as shown in Figure 8 below.

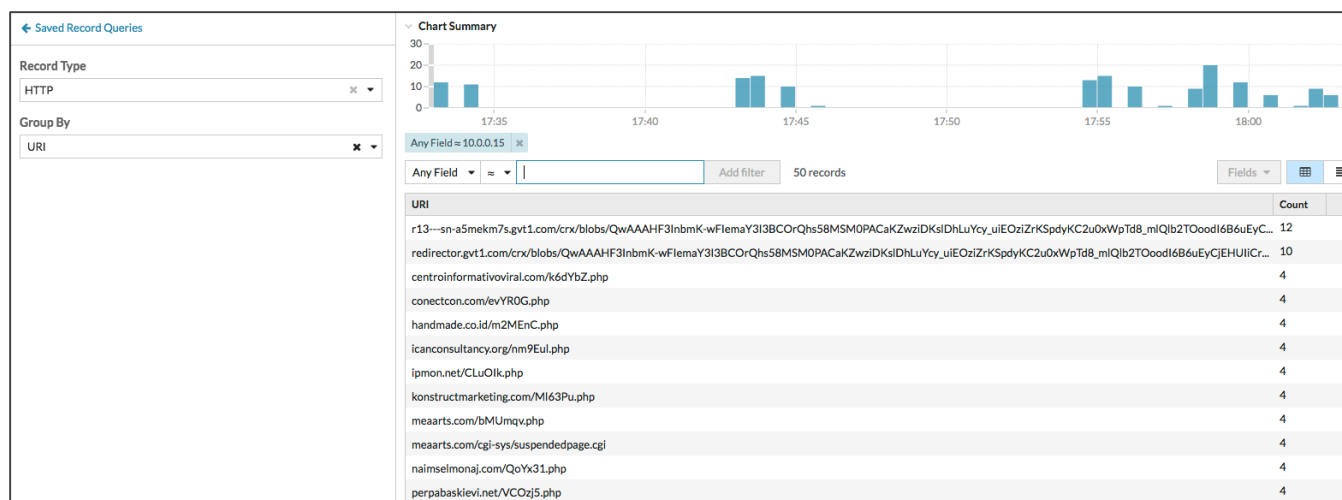


Figure 14: Examining the HTTP transaction records for 10.0.0.15 reveals the source of the malware infection

Finally, three more clicks reveal that a particular suspect URI has only been accessed by this one client across the enterprise, ensuring that personnel and resources are not squandered on an isolated threat.

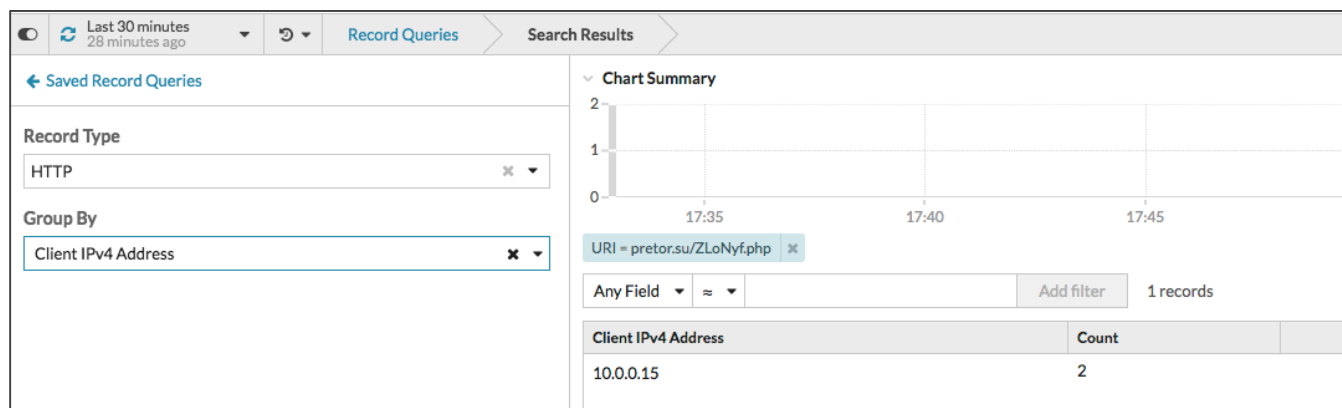


Figure 15: The malicious URI has only been accessed by one client

Hunt Example 4: Russian DNS Queries and DNS Tunneling Detection

In the figure below, a CPT operative is able to interrogate every DNS request made from all clients enterprise-wide in a selected time period and identify which queries are for Russian FQDNs, along with details about the DNS transaction. This output was generated by merely typing “.ru” in the global search field and clicking on DNS Requests from the results in the drop-down menu. Both isolated and DNS beacon behavior can be easily revealed in this manner.

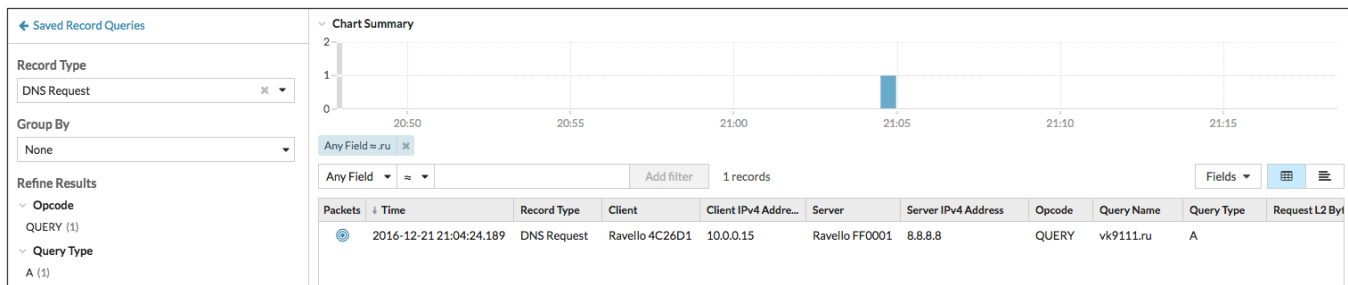


Figure 16: Filtering on “.ru” quickly narrows down DNS request transaction records from Russian domains

Just three more clicks reveal every protocol transaction involving the client that performed this DNS query, allowing a CPT operative to identify all network activity this client has performed, to determine if a malicious payload was downloaded, and to track any subsequent actions made by this payload.

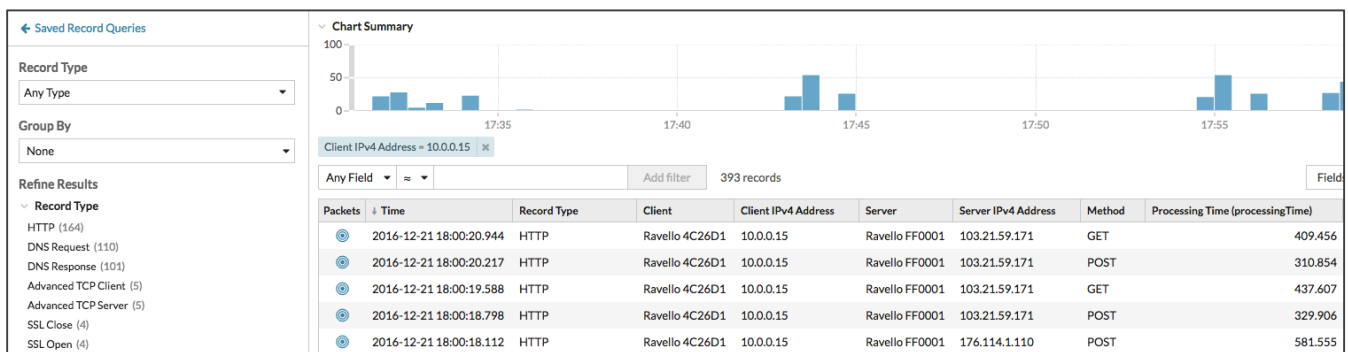


Figure 17: Pivoting the query to focus on the client reveals records for all network activity

Additionally, a glance at a DNS tunneling detection dashboard can reveal whether this activity was part of enterprise-wide anomalous traffic indicative of DNS tunneling. This dashboard tracks multiple characteristics of DNS tunneling behavior, including:

- Unapproved DNS servers
- Large DNS response payloads sizes and high number of DNS responses
- Unexpected use/volume of TXT and NULL records
- Unexpectedly long query/response names

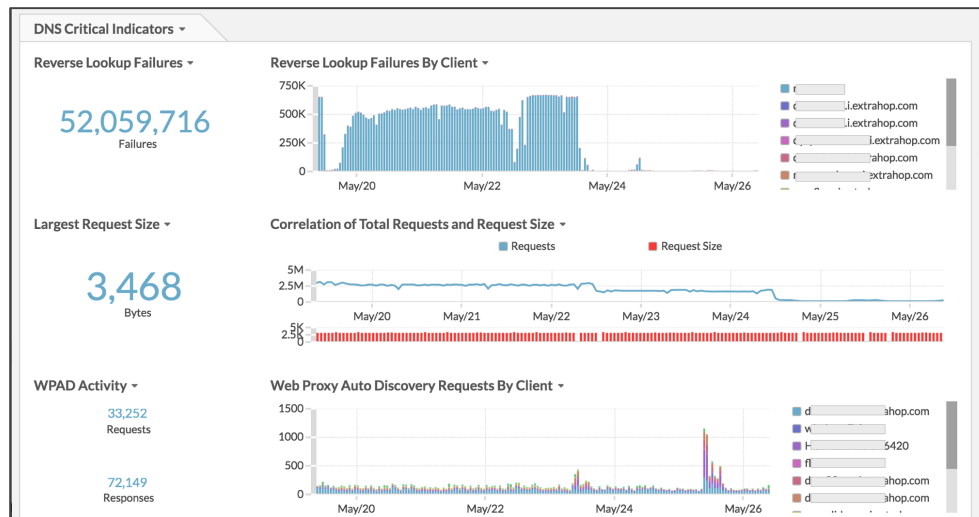


Figure 18: Suspicious DNS activity includes reverse lookup failures, large requests, WPAD activity, and ISATAP tunneling

About the ExtraHop Platform

The ExtraHop platform is a simple turnkey solution that empowers you to make sense of all data passing over the network. Your data in motion is the most valuable source of information that your organization can mine for insights. To access your data in motion, however, you need a platform for transforming large volumes of unstructured network packets into structured wire data. The ExtraHop platform is built to do exactly that at an unprecedented scale.

The ExtraHop platform is a completely passive out-of-band solution, requiring no agents, host configurations, or credentialed access. It will provide maximum visibility into the transactions occurring within your network without degradation or disruption to the environment. It will not actively interrogate any assets, nor will it add any additional traffic to the networks it monitors.




Figure 19: Unstructured data is reassembled into structured wire data that can be mined for insights

Conclusion

Threat hunting is an emerging practice born out of a need to detect more sophisticated threats that evade perimeter defenses and passive monitoring. Early industry feedback is encouraging, with a vast majority (88 percent) of respondents reporting reduced dwell time (the period from initial infection to detection) as a result of their threat hunting efforts.¹

Wire data is an unbiased, real-time source of intelligence but has not been made available to CPTs. The ExtraHop platform unlocks the value of wire data for threat hunting efforts. By using ExtraHop as a real-time threat-hunting platform, you can dramatically

increase the depth and breadth of visibility while decreasing the amount of time and effort needed to derive actionable intelligence. When evaluating platforms for CPTs, it is important to consider the following:

- Does this solution make it easy to collect low-noise, relevant data?
- How easy is it to search through data, derive insight from it, and rapidly act on that insight?
- How easy is integration with existing security workflow and orchestration platforms?
- What impact will this capability have on time-to-detection and time-to-resolution?
- What kind of breadth and depth of information does this solution offer?
- How easy is this platform to deploy and what impact will it have on the environment?
- How susceptible is the monitoring system and data to alteration by malicious actors?

The ExtraHop platform is purpose-built to address all of these considerations, and greatly increases the level of visibility and effectiveness of Cyber Protection Teams.

ABOUT EXTRAHOP

ExtraHop makes real-time data-driven IT operations possible. By harnessing the power of wire data in real time, network, application, security, and business teams make faster, more accurate decisions that optimize performance and minimize risk. Hundreds of organizations, including Fortune 500 companies such as Sony, Lockheed Martin, Microsoft, Adobe, and Google, start with ExtraHop to discover, observe, analyze, and intelligently act on all data in flight on-premises and in the cloud.

ExtraHop Networks, Inc.

520 Pike Street, Suite 1700
Seattle, WA 98101 USA

www.extrahop.com
info@extrahop.com
T 877-333-9872
F 206-274-6393

Customer Support support@extrahop.com
877-333-9872 (US)
+44 (0)845 5199150 (EMEA)