ExtraHop

WHITE PAPER

Reveal(x)

Situational Intelligence for Cyber Hunters

By Sam Richman

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. Unlike reactive forensic investigations, the ideal is to identify and disrupt attack activities before the mission is complete. This white paper explains how Cyber Protection Teams (CPTs) can use high fidelity, real-time network data and analytics to automate detection, speed investigations, and improve the granularity and collection of information. The paper includes two styles: *automated* threat detection, where analytics provide a high-priority starting point for a hunt, and *active* threat hunting, where analytics and data enrich streamlined CPT workflows. Examples show rapid pursuit of brute force attacks, data exfiltration, reconnaissance/lateral movement, ransomware infections, and malicious DNS behavior.

TABLE OF CONTENTS

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 security value and timeliness of self-reported data such as logs, SNMP, and NetFlow metrics.

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.

Who Should Hunt?

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 as separate teams. Organizations should first prioritize securing their infrastructure and building out monitoring capabilities. Next, automated, real-time analysis of the behavior of critical assets can replicate some of the benefits of hunting for organizations that want to integrate these workflows into their security operations. Formal hunt teams can become either a dedicated or a rotational option as maturity increases.

Security practitioners appreciate the idea of seeking out active threats instead of waiting until notified, but few organizations are being as proactive as they would like. In a 2018 survey of 461 cybersecurity professionals, Crowd Research Partners found that respondents spent much more time (60 percent of time) reactively investigating security incidents through activities such as alert triage than they spent proactively seeking out threats (only 40 percent of time). The same survey said that only 24% felt enough time was spent searching for emerging and advanced threats.

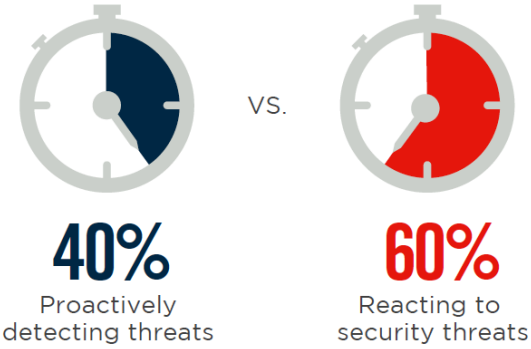

In a typical week, what percentage of your threat management time is spent with alert triage or reactive response to security threats vs. engaging in proactive and innovative detection methods?
Source: The 2018 Threat Hunting Report, Crowd Research Partners

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

Highly accurate, highly fresh data is critical for detecting and disrupting active attack activities. Real-time network traffic analytics can generate authoritative, indexed and complete data to serve as a trusted source of information. Wire data is an observed record of activity, unlike self-reported information such as log, SNMP, and agent data. It is therefore 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 packet capture 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 100 Gbps networks.

[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] 2018 Mandiant M-Trends Report indicating a 101-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

ExtraHop Reveal(x) takes the opposite approach from traditional packet capture, collecting raw network traffic and mining it in real time at 100 Gbps per appliance, automatically discovering client and server assets, and distilling petabytes of traffic per day into manageable and meaningful, structured and indexed data. This transformation 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 50 industry standard L2-L7 protocols to extract transactions and thousands of metrics (also known as dimensions or features) from all these protocols simultaneously.

This real-time approach feeds advanced behavioral analytics, and also allows CPTs to proactively find and explore suspicious activity and behaviors in an efficient and timely manner. Searchable via an elastic engine, expansive data access improves priority, accuracy, and response times between all teams involved in detecting and mitigating an attack. Real-time detections by the ExtraHop Reveal(x) platform allow CPT operators to immediately shift focus to the methods and assets involved in an active attack without being overwhelmed by a huge backlog of data. Without ExtraHop Reveal(x), these insights would require a time-intensive and multi-console 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. Advanced users can customize displays, search for data and pivot through investigations in real time, and set triggers to flag specific activities they want to monitor.

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

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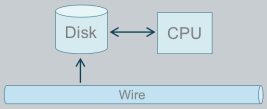	Analyze first, then write to disk 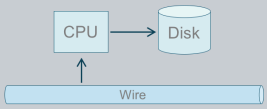
Performance limits	Disk speed	Bus throughput and RAM
Lookback	Data typically stored for days	Data typically stored for months

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) provides CPTs with frictionless access to what matters: a dataset which would otherwise require stitching together of multiple self-reported data sources which, if they even exist, could be compromised by attackers. This affords them the ability to detect, observe, and measure anomalous behavior seen on the network from any device asset or user across all hosts, services, and transactions.

- Automatically discover, classify, and baseline all devices communicating on the network and discover their dependencies
- Identify and investigate anomalies by endpoint, protocol, or user

- Audit critical assets for unauthorized connectivity and protocols/services
- Detect reconnaissance and lateral movement behavior.
- Detect tunneled Command and Control channels and data exfiltration
- Geolocation of traffic sources/destinations

Figure 1: Geolocation of traffic sources and destinations by protocol

ExtraHop Reveal(x) utilizes machine learning to continuously monitor all assets for critical security anomalies. These behavior-based alerts do not require any configuration by your teams. The ExtraHop platform builds baselines for new devices 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. These anomalies serve as effective incident investigation start points, include context to help staff determine the level of severity of the event, and provide paths to guide an operator into the detailed metrics and transactions which characterize the anomaly.

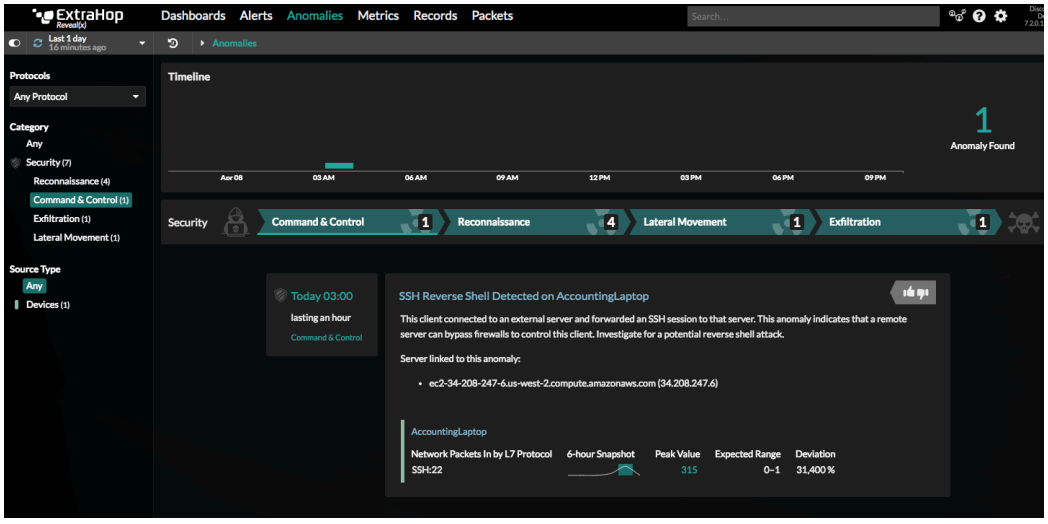

Figure 2: Anomaly detection detects abnormal behavior, such as Command and Control channels

These anomalies can initiate immediate action by REST API integration with security orchestration/workflow platforms like Phantom or ServiceNow, or direct enforcement/mitigation with platforms like Cisco Tetration, Palo Alto, or even custom workflows which support a REST API. For customers who desire specific security indicator detection, JavaScript based triggers are supported, which 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. For 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.

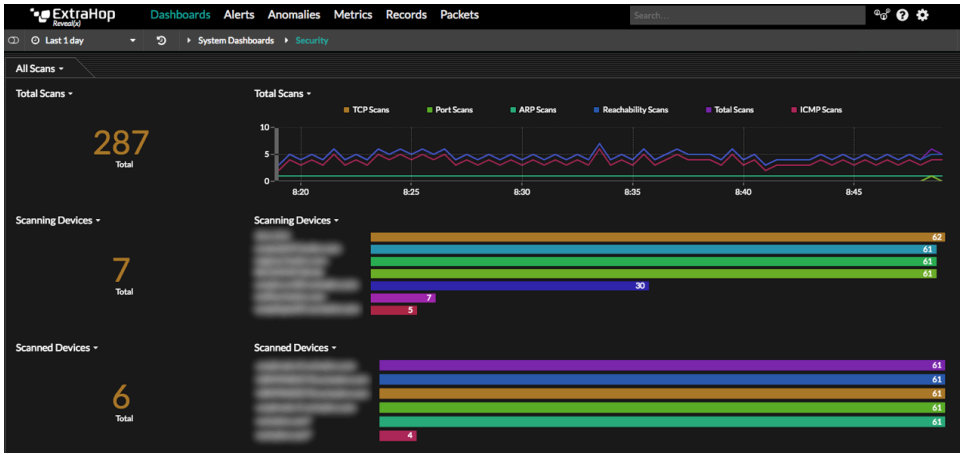

Figure 3: Scan detection

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 devices 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 anomalies, metrics, transactions, and PCAPs within the same intuitive visual user interface, allowing them to achieve a complete incident timeline and forensics record without the need to consolidate data from disparate tools.

Zero Knowledge Discovery of Endpoints 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 effectively implemented and readily audited from this strong security visibility posture. ExtraHop can collect traffic from wherever you desire increased visibility. Typically, East-West traffic is collected from core switches, but you can ingest traffic from the DMZ for North-South visibility, via ERSPAN for intra-VM host traffic, and RPCAPD agents in a cloud deployment.

Cross-Tier Protocol Visibility Using a Visual User Interface

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 automated threat detection and active hunting**.** These hunts leverage the visibility, behavioral analytics, and assisted investigative workflows of ExtraHop Reveal(x).

Hunt Example 1: Brute Force Attack and Data Exfiltration Investigation

A CPT operator is able to easily investigate detected anomalies or hunches by using Reveal(x) 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 red:

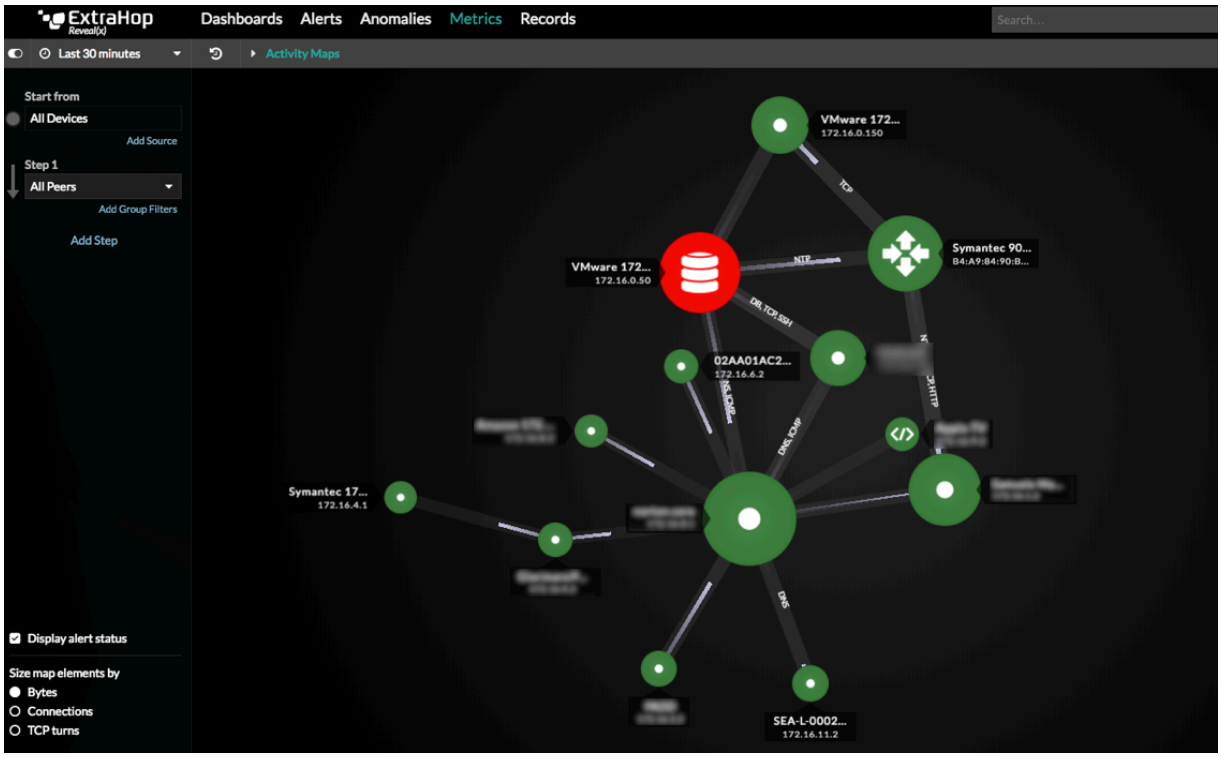

Figure 4: Live Activity Map showing detection of a brute force attack.

One click focuses our attention on the database server under attack, and reveals all other traffic involving the suspect server, which includes outbound network time protocol (NTP) traffic destined for a border router, a violation of security policy. NTP tunneling is a subtle and effective method of data exfiltration[1] because NTP traffic is often ignored when people consider risk in the 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 (figure 6) 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

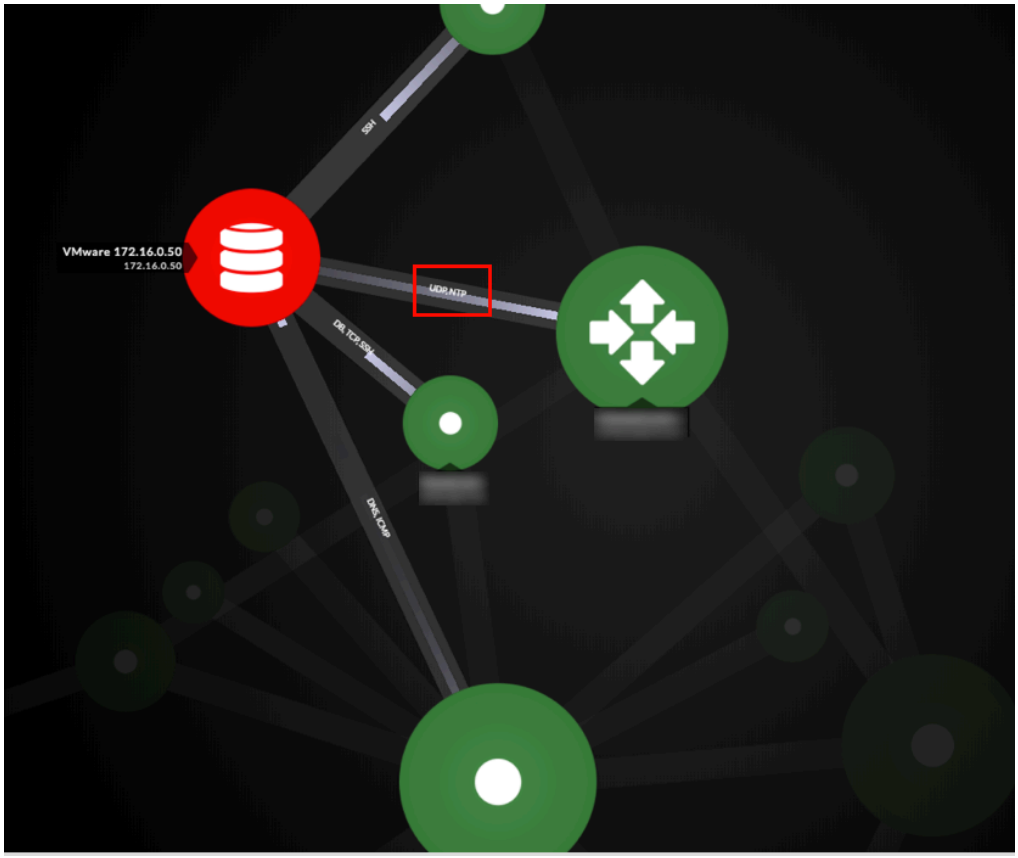

Figure 5: Brute force attack investigation in progress. Egress NTP traffic observed.

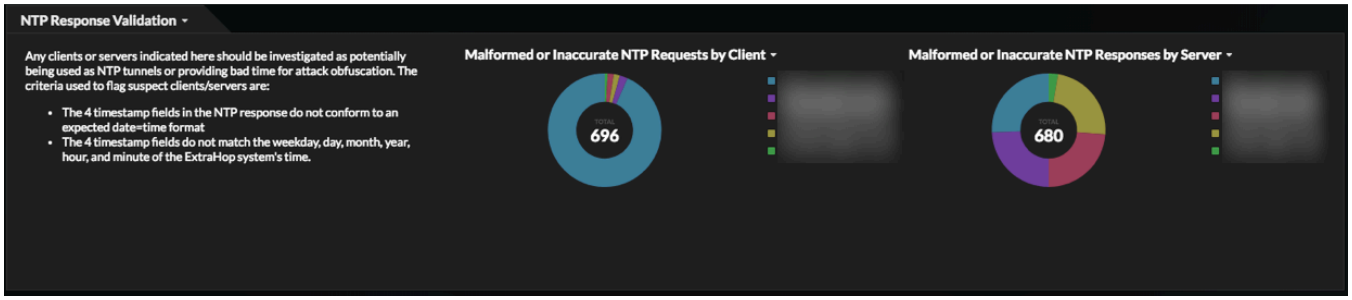

Figure 6: NTP Security Dashboard explains Detections

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.

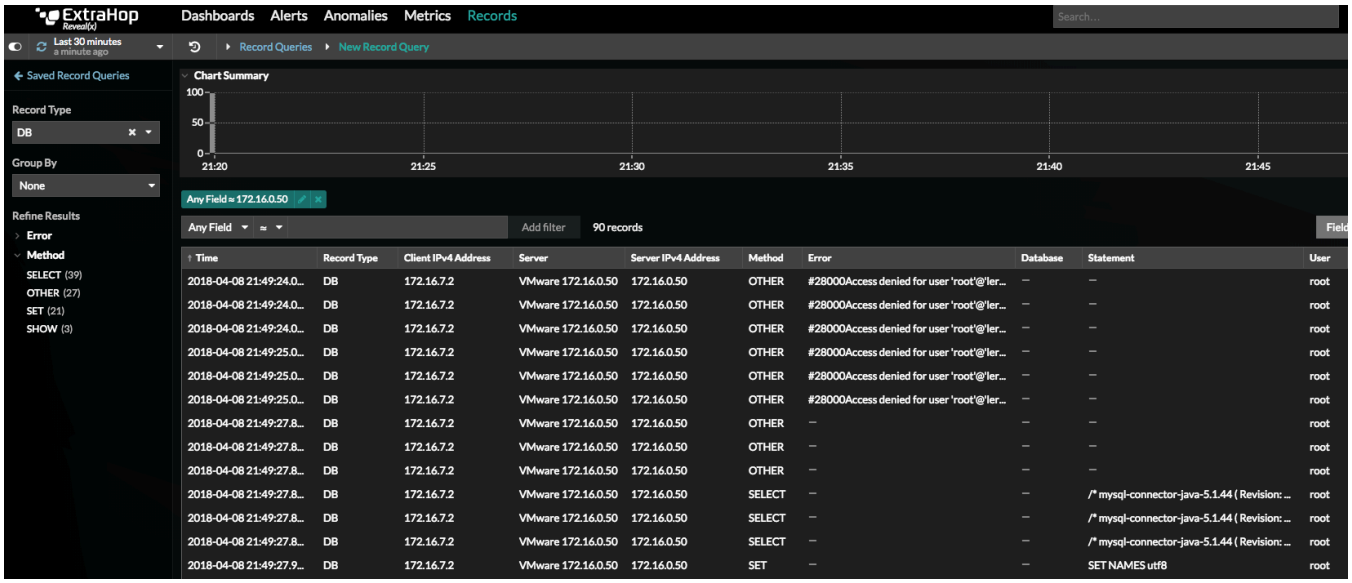

Figure 7: Brute force attack database transaction records

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, Lateral Movement, and Exfiltration 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, getting immediate visibility into the unfolding stages of an attack. Rapid detection of any stage could permit blocking of the attacker before mission completion, but we walk the entire attack to show how the hunt might unfold.

In this example, anomaly detection has revealed evidence of reconnaissance/lateral movement outside of the normal baseline for this enterprise, starting with DNS reverse lookup scans. Since most infrastructures do not log all DNS requests/responses, this reconnaissance technique is difficult to detect.

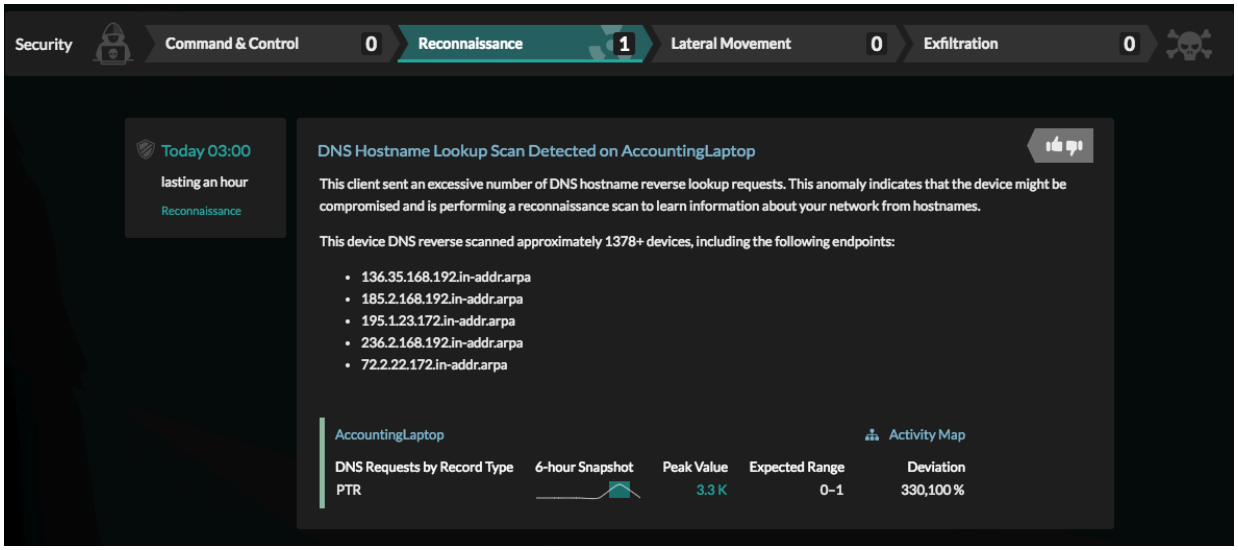

Figure 7: Reconnaissance anomaly detected (DNS reverse lookup scan).

The next contextualized anomalies reveal failed attempts to access network filers and files, additional reconnaissance behavior:

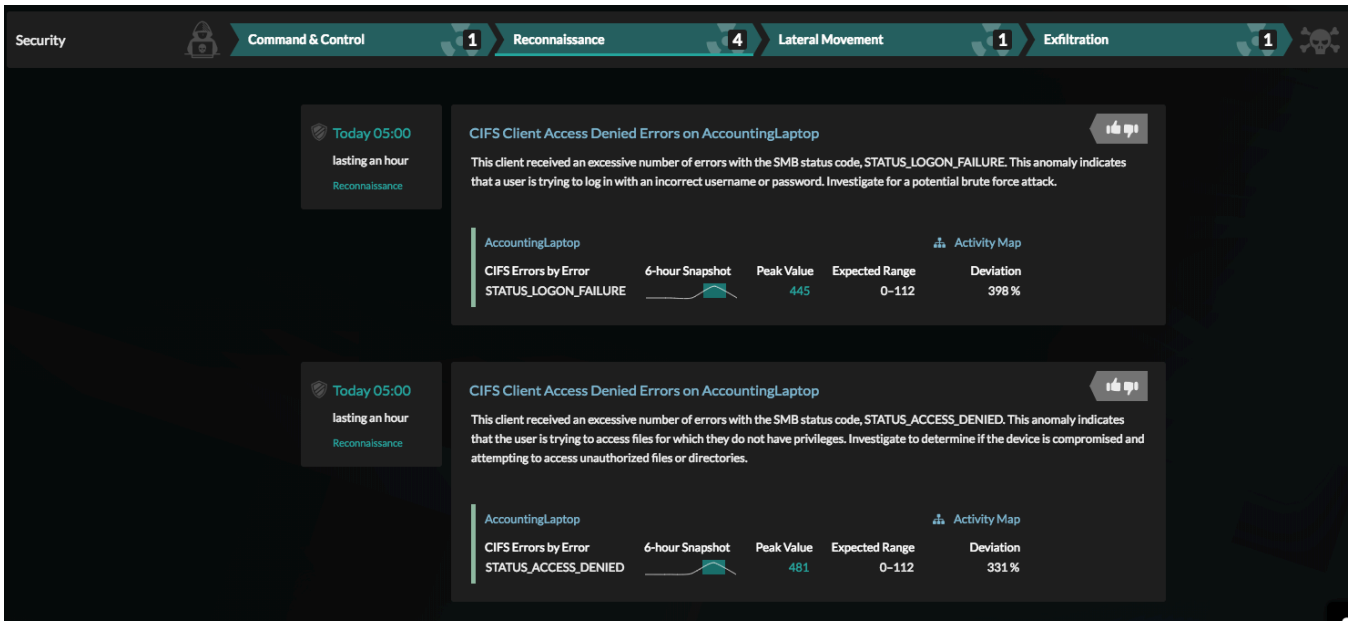

Figure 8: Reconnaissance anomaly detected (network filer access attempts).

The next related anomaly reveals successful file access attempts, an indication of successful lateral movement.

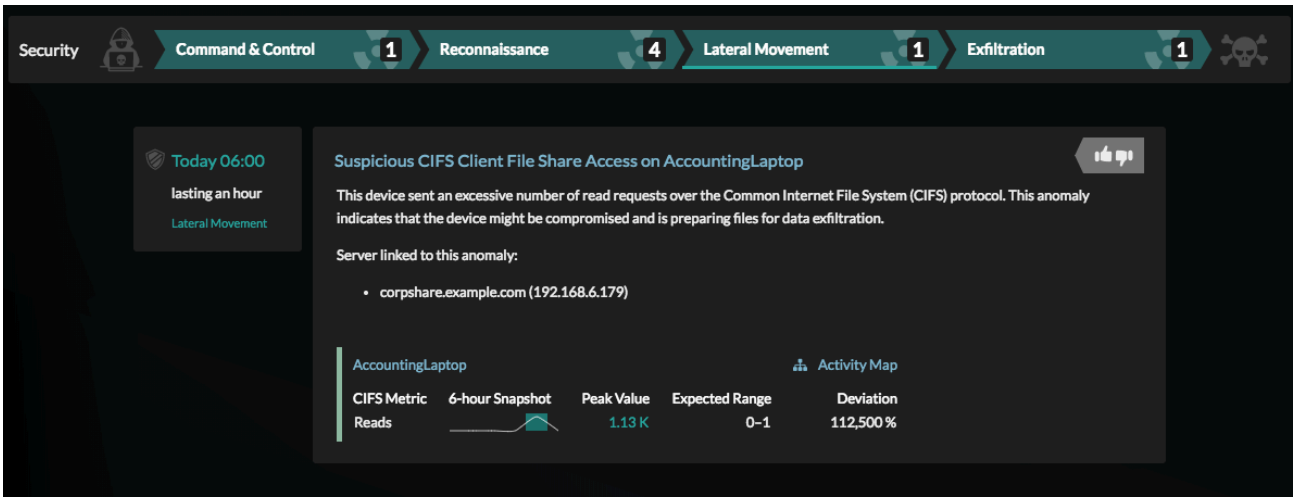

Figure 9: Lateral movement anomaly detected.

Clicking into the lateral movement anomaly reveals a focused Live Activity Map showing the assets involved:

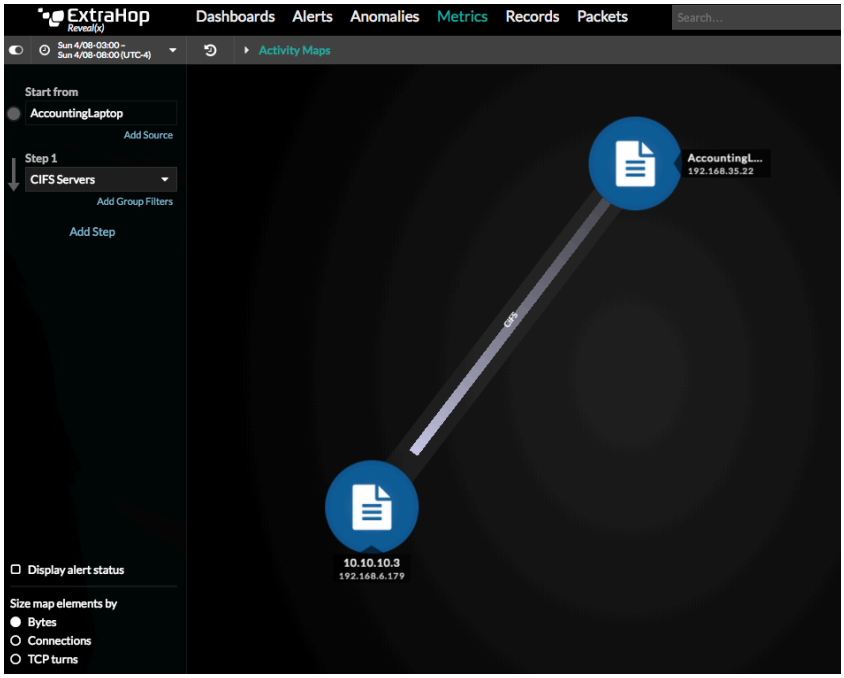

Figure 10: Lateral movement Live Activity Map

Clicking on the connector between the file server and the client reveals the CIFS transactions, showing the precise, sensitive files that are being accessed:

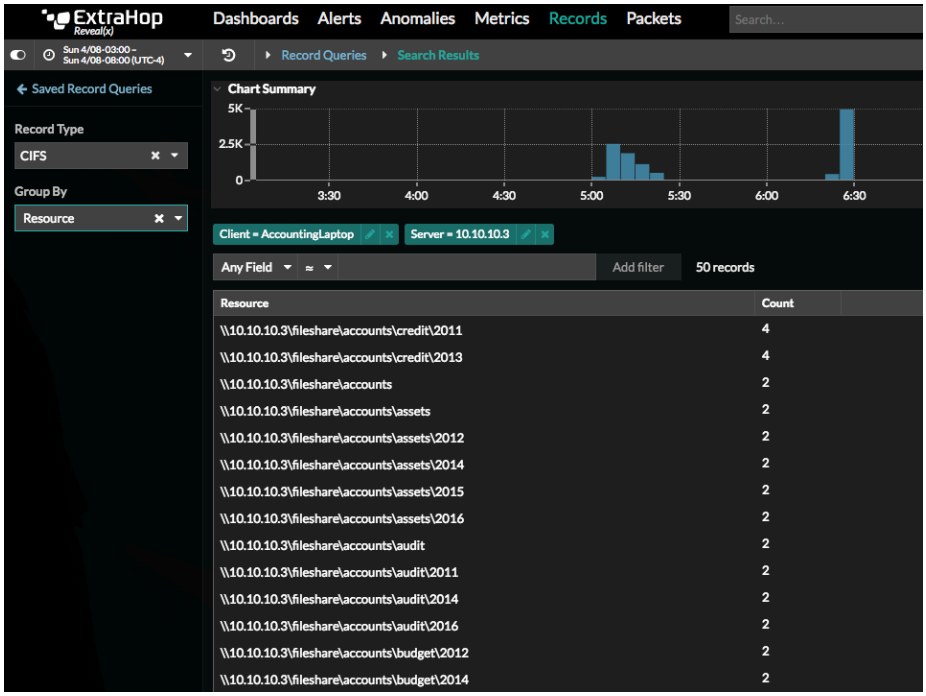

Figure 11: Lateral movement of CIFS transaction records.

A click into the final anomaly in this incident reveals that the asset in question has exfiltrated 1.1 GB of data via SSH. The hunt has gone from DNS anomalies to exploration of CIFS and SSH traffic to fully scope the attack chain within a single investigative environment.

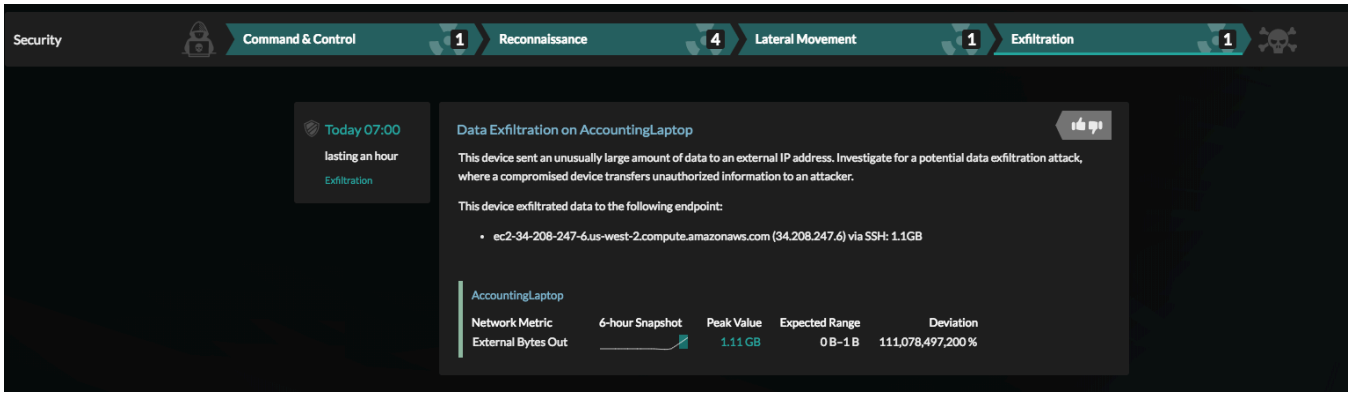

Figure 12: Data Exfiltration anomaly.

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 that initiated the attack. In the following dashboard, clients performing file WRITE/MODIFY operations with suspicious file extensions are revealed. Hunters

can immediately contain this attack by isolating the infected hosts (response should already be planned to include integrations with controls to support this action), and responders can replay the traffic to restore the encrypted files.

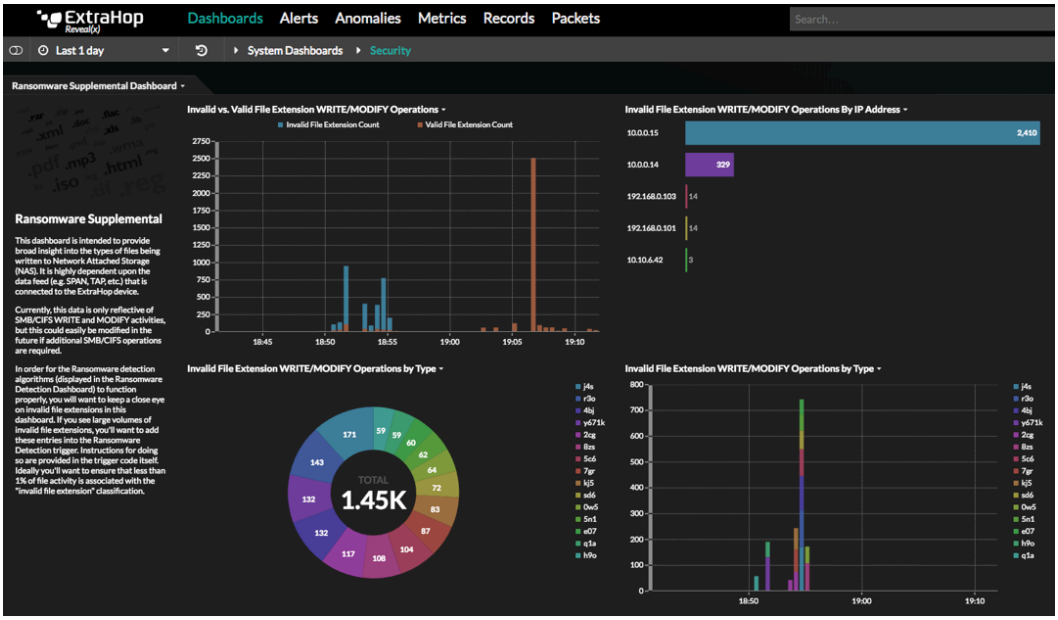

Figure 13: The ransomware detection dashboard reveals infected clients, including 10.0.0.15

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) that reveals the "HELP_DECRYPTING" files that the ransomware package created, shown in Figure 7 below.

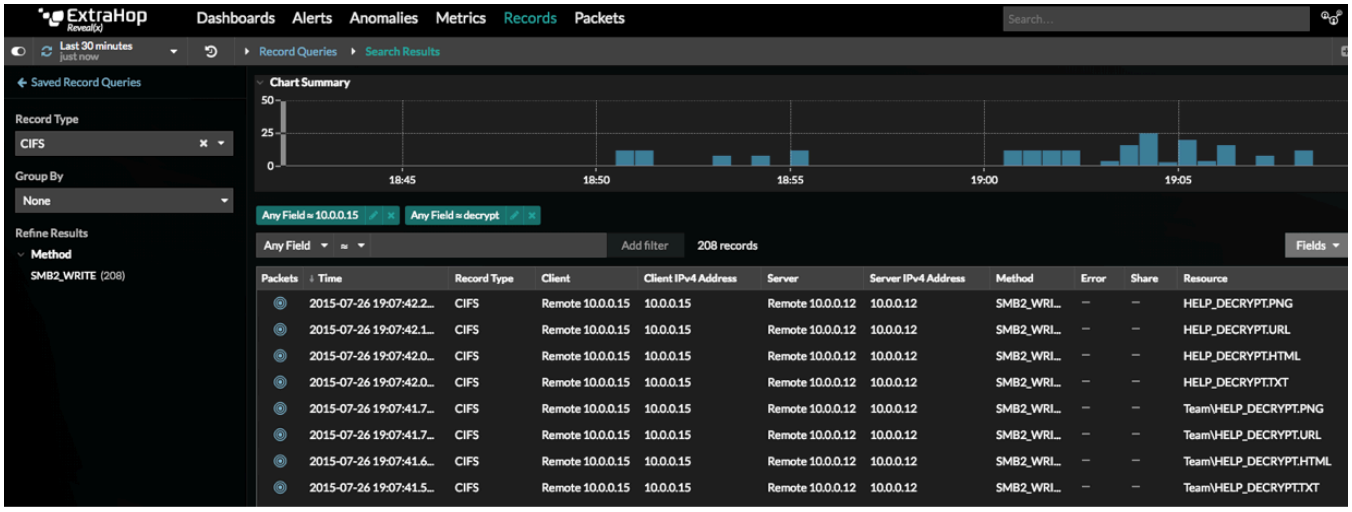

Figure 14: The ransomware package's component files on 10.0.0.15

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.

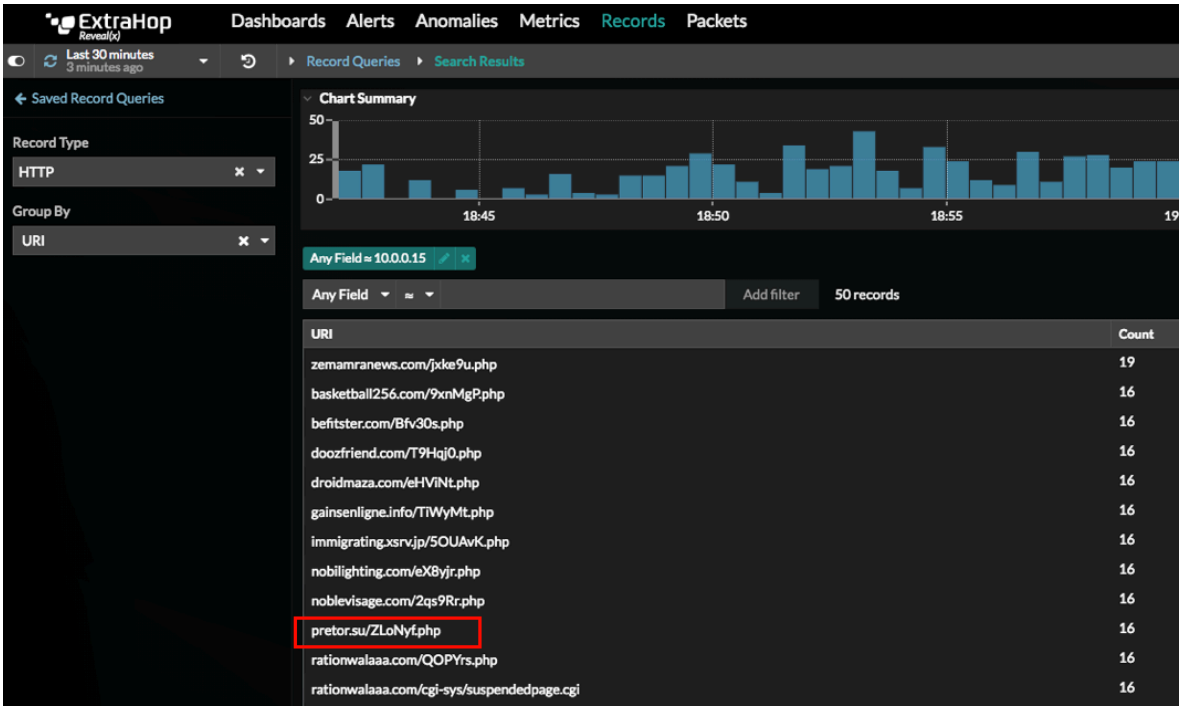

Figure 15: Examining the HTTP transaction records for 10.0.0.15 reveals the source of the malware infection.

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 unnecessary scoping the impact of an isolated threat.

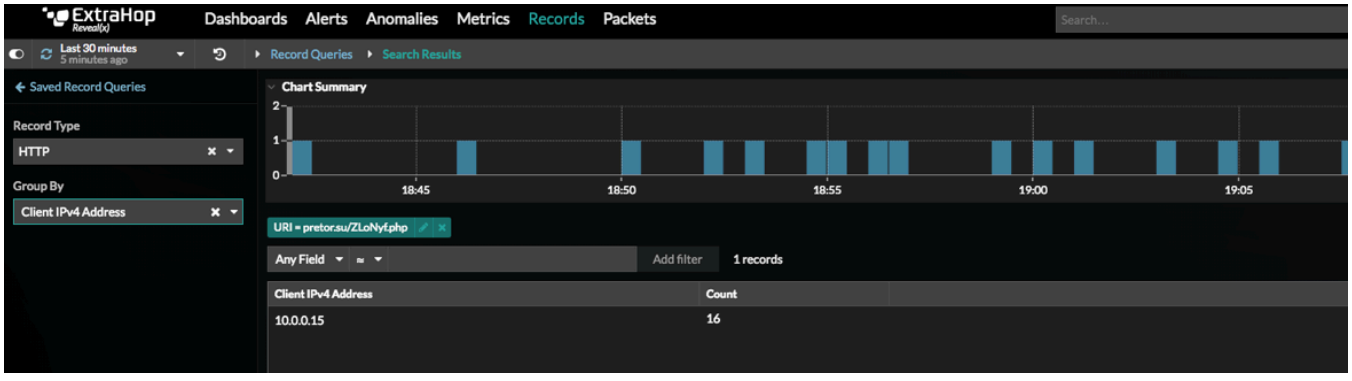

Figure 16: The malicious URI has only been accessed by one client.

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 fully qualified domain names (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, providing action points for containing and investigating the attack.

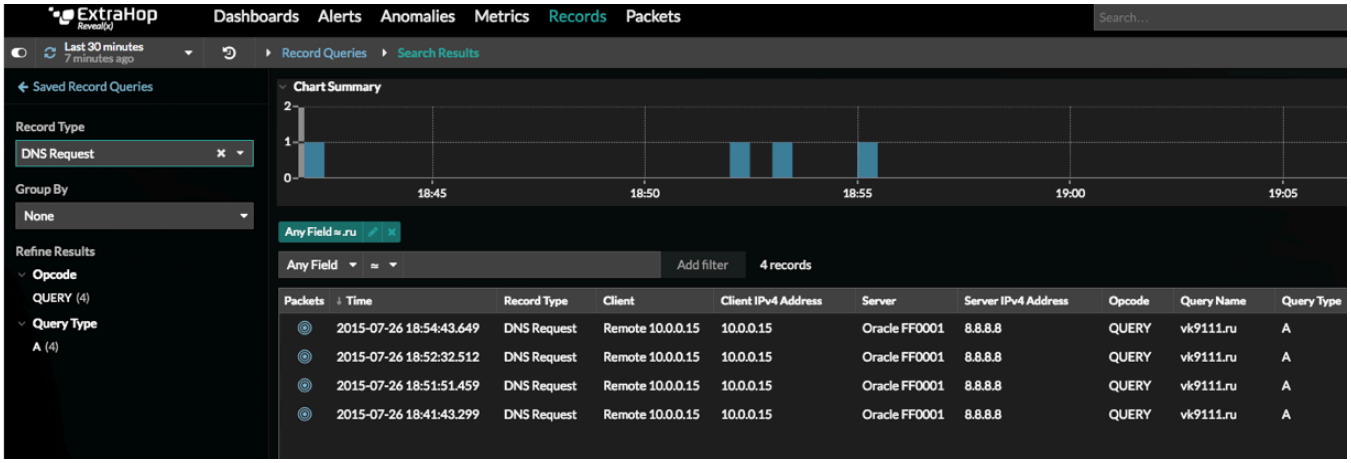

Figure 17: Filtering on ".ru" quickly narrows down DNS request transaction records from Russian domains.

Just three more clicks reveal every protocol transaction involving the client that performed this DNS query, allowing a CPT operator 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.

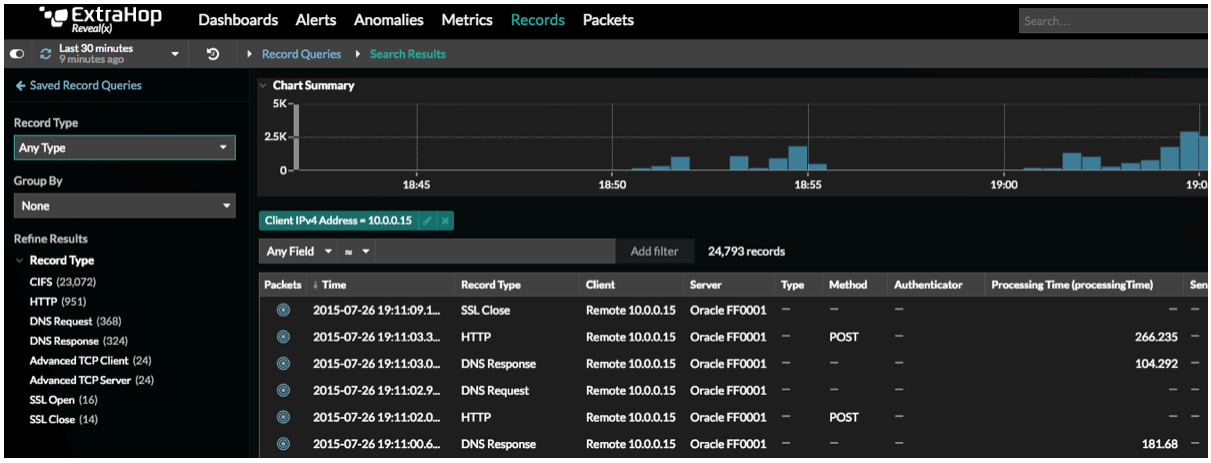

Figure 18: Pivoting the query to focus on the client reveals records for all network activity.

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 or random query/responses

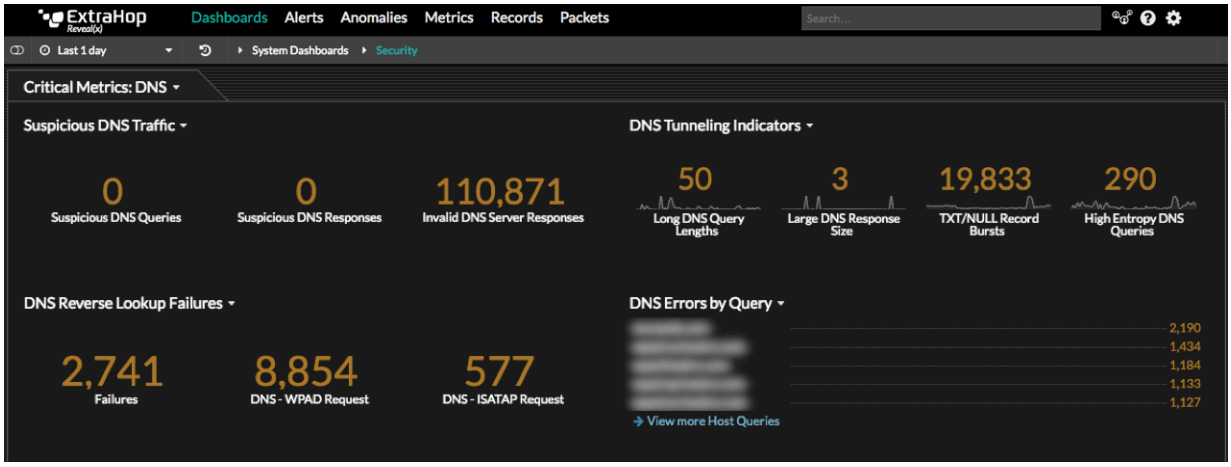

Figure 19: Suspicious DNS activity includes reverse lookup failures, large requests, WPAD activity, and ISATAP tunneling

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 situational intelligence that has not previously been made available to CPTs. The ExtraHop platform unlocks the value of wire data for threat hunting efforts. With ExtraHop automatically discovering, classifying, and monitoring active network assets, the resulting real-time analysis of all transactions on the network (data-in-motion) provides CPTs with frictionless access to what matters: a high-fidelity dataset that would otherwise require stitching together of multiple low-fidelity data sources and manual hunting for evidence. Machine learning performing behavioral analysis on critical assets provides further advantages, approximating advanced hunting expertise and business insights to direct attention to subtle and sophisticated attack activities.

By using ExtraHop as an automated or 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 and surface meaningful anomalies?
- 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 Reveal(x) platform is purpose-built to address all of these considerations, and greatly increases the level of visibility and effectiveness of Cyber Protection Teams.

Appendix: Comparison with Threat Hunting Using Traditional Data Sources

When considering the value of the network traffic analysis that ExtraHop Reveal(x) provides, it is important to discuss the myriad of traditional data sources which would be needed to achieve a comparable level of visibility, along with recognition of their challenges. It is also important to not only consider the time and resources needed to obtain and analyze this traditional data, but also evaluate its reliability during an active attack. As an example of such a comparison, a discussion of the Hunt Examples 1 and 2 follows:

Hunt Examples 1 and 2 IOCs: ***Brute force database attack, unauthorized database access, NTP exfiltration, DNS reverse lookup scans, and network filer reconnaissance/ lateral movement.***

In order to attempt to detect the IOCs that characterize this malicious behavior, the following data sources would be required:

- Database access logs to detect the brute force attack
  - While feasible to collect, a compromise of the database server would likely involve obfuscation or elimination of these logs.
- Database audit logs to detect the actions taken by the threat actor
  - Logging all database transactions has heavy resource requirements and is also susceptible to compromise by a threat actor
- Egress firewall logs to detect the rogue NTP traffic
  - Logging all outbound traffic, while possible, requires significant resources, and is also unlikely to reveal the NTP traffic as an exfiltration channel.
- DNS query/response logs
  - With almost every application transaction involving at least one DNS query/response, this is one of the chattiest datacenter protocols, and requires significant resources to collect and store. Additionally, rogue/compromised and misconfigured DNS servers will certainly not send logs.
- Network filer access logs
  - High-granularity access logs are impractical to collect and store, and while they may be reliable, the sheer volume can hide malicious behavior.

The challenges discussed above can also be considered in the context of three fundamental characteristics of any security data source:

- Timeliness: Since log data is self-reported, log aggregators must first wait until assets transmit their telemetry and are thus at the mercy of host and network-based delays and failures. Once received, this telemetry must still be indexed, correlated, and analyzed before it can be effectively used to understand a security incident. These factors significantly affect the speed at which security events can be acted upon. The wide variety of logs involved in these hunt examples exacerbates this challenge.
- Accuracy: Logs and other self-reported data vary dramatically in both format and included information, and commonly lack consistency. Because of this, log aggregators require significant configuration in order to accommodate this inconsistency. Even then, the results of analysis are still the systems' interpretation of the data being received, interpretations which when inaccurate, give rise to false positives and negatives.
- Reliability: Because log aggregators rely on self-reported data, the implication is that an asset must be known about, configured to send this data, and most importantly, not have been intentionally misconfigured to hide an attacker's

actions. The last point is critical: it is common practice for threat actors to reconfigure compromised assets to send intentionally misleading or inaccurate data to obfuscate malicious activity.

ExtraHop Reveal(x) avoids the challenges above by providing agile and real-time access to trustworthy security data directly from the network and includes automatic asset discovery/classification to continuously validate the completeness of one's understanding of the infrastructure. This strategy ensures the objective and consistent reporting of asset behavior, regardless of the asset's ability to self-report data.

ABOUT EXTRAHOP

ExtraHop is the first place IT turns for insights that transform and secure the digital enterprise. By applying real-time analytics and machine learning to all digital interactions on the network, ExtraHop delivers instant and accurate insights that help IT improve security, performance, and the digital experience. Just ask the hundreds of global ExtraHop customers, including Sony, Lockheed Martin, Microsoft, Adobe, and Google.

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)