



Post-Intrusion Report

June 2015

TABLE OF CONTENTS

About the data in this report.....	3
Classification of data.....	3
High-level trends.....	4
Detected threats by category.....	4
A spike in indicators of targeted attacks.....	5
A view by industry.....	5
Focus on command and control.....	6
Hidden tunnels.....	7
Focus on botnet monetization.....	7
Focus on reconnaissance.....	8
Focus on lateral movement.....	9
Focus on data exfiltration.....	9
Conclusions.....	10

The Vectra® Networks Post-Intrusion Report (PIR) provides a first-hand analysis of active and persistent network threats inside an organization. This study takes a multidisciplinary approach that spans all strategic phases of a cyber attack, and as a result reveals trend related to malware behavior, attacker communication techniques, internal reconnaissance, lateral movement, and data exfiltration.

Key Findings:

- 100% of the networks analyzed in the report exhibited one or more indicators of a targeted attack.
- Targeted attack indicators were on the rise, led by a 580% increase in lateral movement techniques along with a 270% increase in internal reconnaissance. A spike in these behaviors may indicate that attackers are increasingly successful at penetrating perimeter defenses.
- While command and control behaviors remained flat, the riskiest forms of command and control were on the rise with a marked increase in Tor as well as external remote access tools.
- For the first time, Vectra was able to perform a study of hidden tunnels without the need to decrypt SSL. This analysis showed that HTTPS is the preferred vehicle over HTTP for hidden tunnels.

About the data in this report

The data in this report is based on metadata from Vectra customers and prospects who opted to share detection metrics from their production network environments. Vectra software identifies active threats by directly monitoring network traffic on the wire in these environments. Vectra monitors both internal host-to-host traffic as well as traffic to and from the Internet to ensure full visibility and context of all phases of an attack. As a result, the data in this report offers a first-hand analysis of active, “in situ” network threats that bypassed next-generation firewalls, intrusion prevention systems, malware sandboxes, host-based security solutions, and other enterprise defenses.

Classification of data

Vectra software automatically detects a variety of threats and attack techniques throughout all phases of a persistent attack. These detections are categorized based on the following strategic phases of an attack

- Command-and-control communications
- Botnet monetization
- Internal reconnaissance
- Lateral movement
- Data exfiltration

The diagram below provides a high-level overview of how each phase relates to one another in the context of a persistent network attack. However, it is important to point that real-world attacks will not necessarily follow this model sequentially or exhibit every phase of an attack.

The phases of attack are further classified into categories based on whether the attack appears to be targeted or opportunistic. In targeted attacks, the attackers seek out data or assets that are unique to the victim network or organization. These assets can include customer, partner or employee data, financial information, and a variety of intellectual property and trade secrets.

Conversely, opportunistic attacks seek out resources and are largely unconcerned with the specific assets of the victim organization. Botnets are a common example of this category.

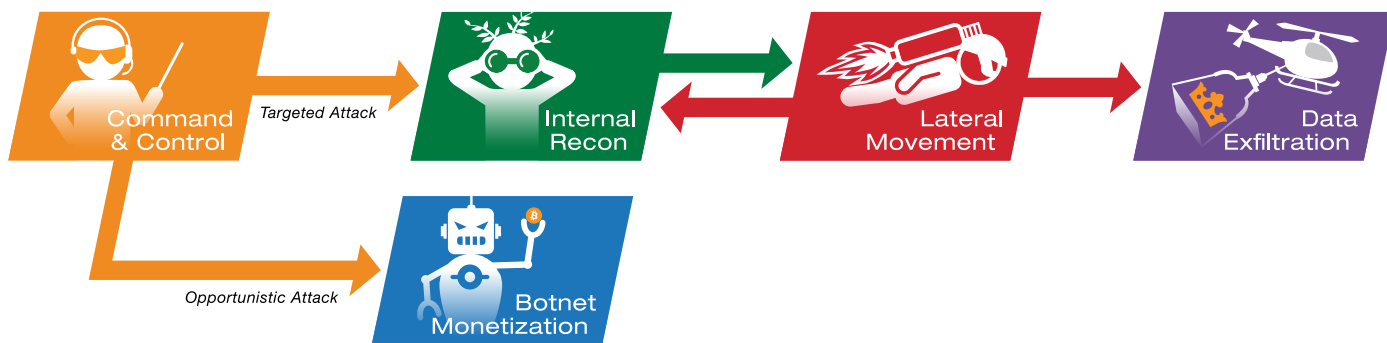


Figure 1: The phases of targeted and opportunistic attacks

For example, a spamming botnets often aim to grow to the largest size possible and infect as many host devices as possible. In these cases, an attacker may use hosts within an organization to send spam messages, but is not concerned with stealing company data.

While these categories can be useful in tracking the apparent intent of an attacker, it is important to note that they are neither mutually exclusive nor permanent. Even basic botnet malware will often attempt to steal user credentials, which could be used in more targeted phases of an attack at a later time.

High-level trends

In this edition of the Vectra Post-Intrusion Report, we examine techniques and behaviors of real-world cyber threats, and for the first time track how these behaviors change over time. The first version of this report was released in November of 2014, and this edition illustrates some of the changes we observed in real-world attacks.

It's important to review the sample size of the data in this report as well as the demographics of participating organizations. This study includes data from 40 organizations spanning a variety of industries including education, energy, engineering, financial services, government, healthcare, legal, media, retail, services, and technology.

These networks contained 248,198 hosts, representing a 122% increase compared to the 111,589 hosts seen in the previous report. Accordingly, the overall number of detections also rose, although at a slightly smaller rate. Vectra observed 46,610 total threat detections, which represents a 97% increase compared to November 2014.

Detected threats by category

Vectra classifies threats into five categories based on the lifecycle of a persistent attack. These categories are (1) botnet monetization, (2) command and control, (3) internal reconnaissance, (4) lateral movement, and (5) data exfiltration. The table below summarizes the relative volume of detections in each category.

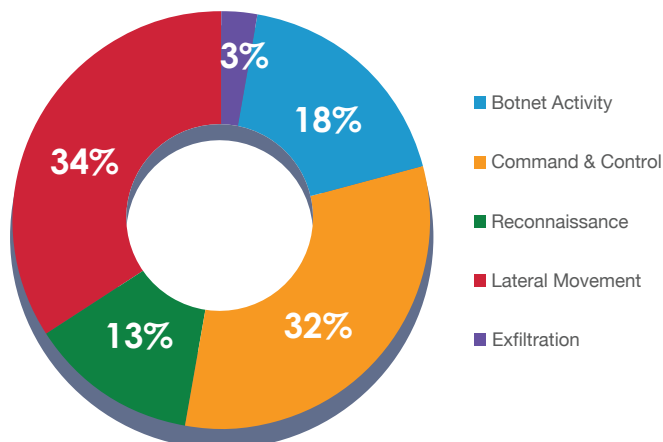


Figure 2: Summary of all detections by detection category

Lateral-movement detections, which track the internal spread of malware and authentication-based attacks such as the use of stolen passwords, led the pack with over 34% of total detections.

- Command and control detections, which identify a wide range of malicious communication techniques, were close behind with 32% of detections.
- Botnet monetization detections track the various ways criminals make money from ad click-fraud, spamming behavior, and distributed denial of service (DDoS) attacks. These botnet-related behaviors accounted for 18% of all detections.
- The reconnaissance category looks for internal reconnaissance performed by an attacker already inside the network and represented 13% of detections.
- Exfiltration detections look for the actual theft of data. The good news here is that it was by far the least common category of detection at 3%.

These results mark a significant shift compared to the previous report, where command and control detections dominated, accounting for 60% of all detections. This shift was driven by significant increases in lateral movement and reconnaissance behaviors, while the growth of command and control remained relatively flat.

The chart below puts these trends into perspective by comparing the percent growth of detections between the last report and this one. While the number of detections nearly doubled, lateral movement and reconnaissance significantly outpaced this benchmark with 580% and 270% growth, respectively.

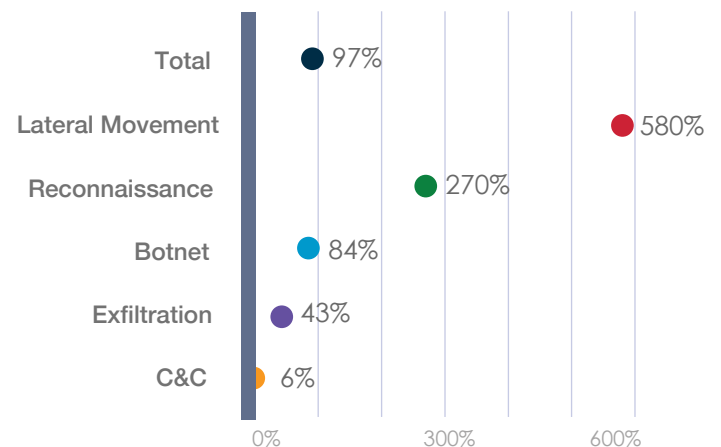


Figure 3: The percentage growth in number of detections overall and by category. Compares data from Q2 CY 2014 to Q2 CY 2015

We examine each of these categories in detail later in the report, but to give a preview, both Lateral Movement and Reconnaissance detections were up across the board and were not tied to one specific type of detection. However, there were industry-specific trends within particular detections.

A spike in indicators of targeted attacks

As mentioned in the Classification of data section at the beginning of the report, Vectra distinguishes between opportunistic and targeted threats whenever possible. Targeted attacks almost always represent the greatest risk to an organization because they can expose customer data, financial information, intellectual property, and trade secrets.

While attacks against retailers like Target and Home Depot have been well documented, the past six months has shown that all industries are at risk. Multiple major breaches against healthcare giants Anthem and Blue Cross exposed a variety of patient data, and the attack against Sony Pictures exposed executive communications, leaked unreleased movies, and created an immense amount of damage.

In the previous report, Vectra observed indicators of a targeted attack in every network that was monitored. This is a stark reminder that threats, whether human or driven by malware, are consistently able to penetrate traditional defenses and take hold in a target network. The marked increase in lateral movement and reconnaissance behaviors is particularly significant because these attack phases are strategic to the success of a targeted attack. These attacks require attackers to persist within a network and spread through the environment.

Consequently, detecting the presence of internal reconnaissance and lateral movement represents one of the most important opportunities to mitigate these threats before assets are compromised.

A view by industry

To dig a bit deeper into threats seen in real networks, Vectra is for the first time providing a breakdown of detection statistics by industry. Figure 4 below shows the percentage of threat detections that were triggered in each industry.

For example, education accounted for 8% all detections analyzed over the course of the study. The columns to the right show the percentage of detections for each detection category.

For each industry, has highlighted in red the threat categories that were most common for that particular industry. For example, lateral movement detections were most common in education, while exfiltration detections were most common in the energy sector.

Figure 4 shows that the technology sector accounted for the most detected threats overall as well as the most threat categories. This indicates that technology companies are well represented in the sample set and contain some of the largest networks.

To account for this sample bias, Vectra also performed an analysis of the number of detections per host seen on the network. This view provides an entirely different perspective by showing how each industry fared per capita as well as which industries generated the most detections by volume.

Industries	Total Detections	Botnet Detections	Command & Control Detections	Exfiltration Detections	Lateral Movement Detections	Reconnaissance Detections
All Industries	100%	18%	32%	3%	34%	13%
Education	8%	9%	4%	4%	12%	4%
Energy	8%	8%	11%	12%	8%	4%
Engineering	3%	2%	2%	2%	5%	3%
Financial	4%	6%	1%	7%	5%	5%
Government	13%	13%	8%	2%	20%	6%
Healthcare	14%	9%	19%	11%	16%	8%
Media	11%	16%	11%	34%	5%	12%
Services	2%	3%	1%	4%	3%	2%
Technology	37%	33%	43%	25%	27%	57%

Figure 4: Percentage of threat detections by type and industry

Detections per 100 hosts Industries	Total	Botnet	Command & Control	Exfiltration	Lateral Movement	Reconnaissance
All Industries	18.78	3.39	6.08	0.51	6.32	2.48
Education	11.15	2.39	1.85	0.18	6.05	0.68
Energy	51.20	9.41	19.40	2.13	17.41	2.86
Engineering	48.51	6.71	8.41	0.94	26.82	5.62
Financial	17.62	4.61	1.93	0.86	7.18	3.04
Government	37.95	7.38	7.29	0.14	20.87	2.27
Healthcare	21.38	2.60	8.57	0.45	8.23	1.54
Media	17.15	4.98	5.11	1.54	2.99	2.54
Services	27.48	6.28	4.81	1.30	12.28	2.80
Technology	13.51	2.32	4.73	0.26	3.49	2.72

Figure 5: Count of threats per 100 hosts by type and industry

Figure 5 shows the results in terms of the number of detections per 100 network hosts. The average number across all industries is shown in gray. Red denotes above average detections per host while green is below average. While technology generated a large volume of detections, the sector fared quite well on a per-host basis.

Focus on command and control

Command-and-control traffic is a key enabler for virtually all types of modern cyber attacks and threats. Virtually any threat that leverages malware will want to control that malware and retrieve information from it.

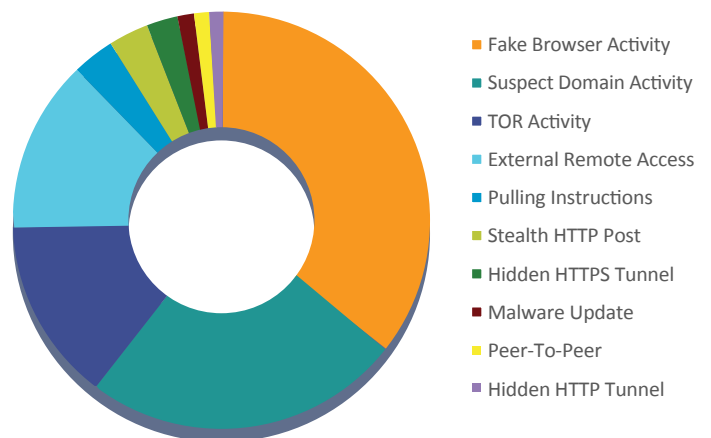
This makes command-and-control, or C&C, communications vital to attackers and a key target for security vendors. This has led to a constantly evolving cat-and-mouse game between security vendors and attackers.

The command-and-control behaviors observed in this report particularly interesting because they represent communications that are not being controlled by firewalls, IPS systems, malware sandboxes or other traditional methods.

This time, the data showed that some old friends are still alive and well, but also marked a change to a variety of evolving techniques used by attackers. As seen in the previous report, the use of fake browsers and newlygenerated domains were the most common ways that attackers tried to hide their traffic.

Fake browser activity allows an attacker to blend in with the massive flows of web traffic seen on most networks. Conversely, suspect domain activity shows that, despite URL blocking rules and faster and faster data feeds, attackers continue to create and use new URLs faster than signatures can be written and delivered.

While the No. 1 and No. 2 techniques within command and control remained steady, there were major changes in the rest of the field. Tor introduced the biggest change, posting a more than 1000% increase and accounting for 14% of all command-and-control traffic.



Fake Browser Activity	5373	36%
Suspect Domain Activity	3715	25%
TOR Activity	2124	14%
External Remote Access	2008	13%
Pulling Instructions	510	3%
Stealth HTTP Post	468	3%
Hidden HTTPS Tunnel	366	2%
Malware Update	193	1%
Peer-to-Peer	175	1%
Hidden HTTP Tunnel	157	1%

Figure 6: Comparison of observed command-and-control behaviors

TOR provides anonymity across the Internet and makes it virtually impossible to track where traffic is going to or coming from. Malware authors and attackers have been taking note of this advantage and using Tor more and more as part of their attack infrastructure.

External remote access also made a considerable jump. These techniques took over the No. 4 spot on the list and posted a 183% increase compared to the previous report.

External remote access is the Vectra name for a variety of human-enabled remote access tools, or RATs. They're indispensable to remote targeted attacks because they use an internal infected host to directly view and control the attack. This hands-on approach is key to most sophisticated attacks, and being able to detect these techniques is critical to defending a network.

Hidden tunnels

Attackers always look for new ways to hide their traffic. One of the most successful methods involves tunneling their traffic within another allowed protocol. With tunneled behavior, the attacker's communication is so customized that it's almost impossible to detect the internal protocol via traditional means.

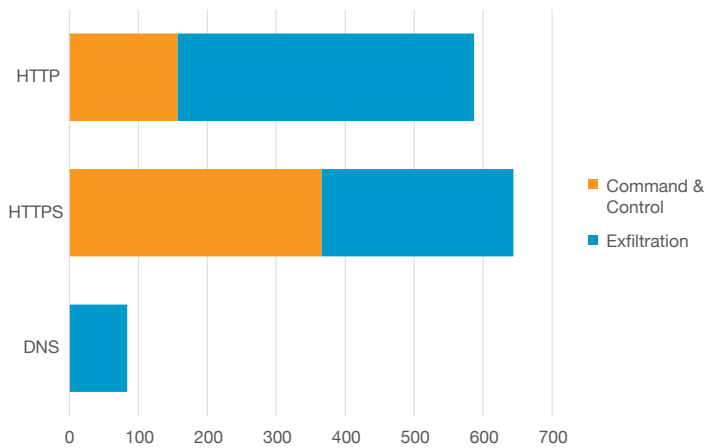
For example, the attacker may use a normal benign HTTP communication as a vehicle, but embed coded messages in text fields, headers or any number of parameters within the session. By riding within such an allowed protocol, the attacker can communicate back and forth without detection.

Vectra software is uniquely suited to identify this type of evasion. Unlike an IPS or next-generation firewall that may attempt to decode the protocol, Vectra applies data science to the communication pattern itself.

If an HTTP session is carrying a hidden second conversation, there will be discernible patterns in the timing, volume and sequencing of traffic. By learning these patterns, Vectra has built models to identify hidden tunnels within HTTP, HTTPS and DNS.

More importantly, this approach does not require direct visibility into the payload of the traffic. This means Vectra can detect hidden tunnels with HTTPS without decrypting the traffic. As a result, Vectra can for the first time to perform a direct and equal comparison of hidden tunnels in both encrypted and clear traffic.

In addition to building data science models for these three protocols, Vectra scientists distinguish between the use of these techniques for command and control as opposed to exfiltration. Obviously, a hidden pipeline of information is useful to an attacker in both contexts.



	DNS	HTTPS	HTTP
Command & Control	0	366	157
Exfiltration	84	278	430

Figure 7: The prevalence of hidden tunnels within HTTP, HTTPS, and DNS

Leveraging this unique insight into hidden tunnels, Vectra researchers were able to directly compare the prevalence of hidden tunnels within HTTP, HTTPS and DNS. This type of analysis has never been practical using traditional methods because it required the decryption of SSL across the entire data set.

By analyzing hidden tunnels without the need for decryption, Vectra researchers were able to perform an apples-to-apples comparison and observed that HTTPS was the most common protocol used for hidden tunnels.

Within command-and control-traffic, HTTPS was over twice as common compared to HTTP. Given that HTTP is typically more common than HTTPS in regular Internet* and enterprise network traffic**, this shows that attackers prefer using HTTPS as an added layer of protection for their communications

Focus on botnet monetization

Botnets are typically large-scale networks in which a centralized criminal or group of criminals infects and coordinates large numbers of hosts to a common goal. A criminal can use this massive network of infected hosts to do any number of things, including send spam messages or drive massive click-fraud campaigns.

Unlike targeted attacks, most botnets are unconcerned with who they infect, but rather are simply trying to infect as many nodes as possible. In most cases, an attacker views the hosts in his botnet as a commodity that performs some rote task on his behalf.

Vectra generically classifies these types of attacks as opportunistic, simply to distinguish them from threats in which an attacker specifically seeks an organization's assets.

* Sandvine Global Internet Phenomena Report 2014

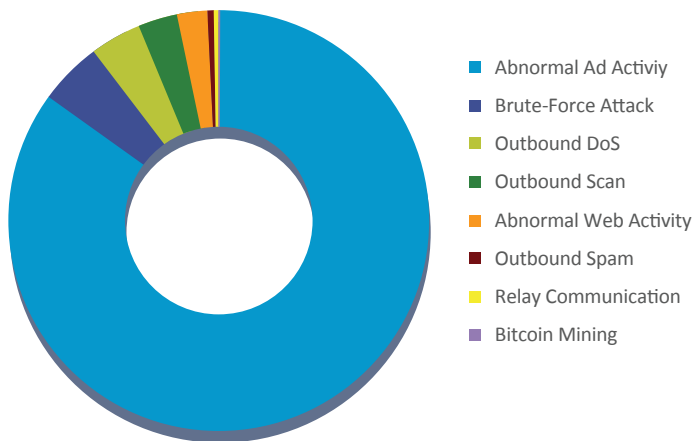
** Palo Alto Networks Application Usage and Threat Report 2014

This does NOT mean that botnets and botnet monetization behaviors pose no risk to an organization. At a minimum, a host that is a node in a botnet is a machine in your environment that is infected with malware and under the control of an outside actor.

An infection that is simply mining bitcoins one day can easily turn to something more serious. Additionally, a botnet could significantly impact the reputation of your network and organization if your network is observed emitting spam or DDoS traffic.

Overall botnet monetization behaviors grew linearly compared to last year's report. And again, ad clickfraud was far and away the most commonly observed botnet monetization behavior, accounting for 85% of all botnet detections.

While not on the scale of the click-fraud behavior, there were some more sinister botnet behaviors. Brute force attacks and outbound DoS attacks were the second and third most common behaviors, respectively. Botnets are common platforms for testing and breaking passwords, which can then be used in more targeted attacks.



Behavior	Count	Percentage
Abnormal Ad Activity	7156	85.0%
Brute-Force Attack	412	4.9%
Outbound DoS	323	3.8%
Outbound Scan	261	3.1%
Abnormal Web Activity	190	2.3%
Outbound Spam	45	0.5%
Relay Communication	27	0.3%
Bitcoin Mining	7	0.1%

Figure 8: Comparison of observed botnet monetization behaviors

These botnets will often test common passwords or username-password combinations stolen in a previous attack in order to find vulnerable targets. By using a large number of infected hosts from many locations, attackers are able to avoid rate-limiting controls on the target site.

DoS attacks were observed emanating largely from networks in the education sector, while brute-force attacks were more evenly distributed across industries.

The Vectra analysis shows that HTTPS is the most popular carrier of hidden tunnels because encryption provides an additional layer of protection to hide the attacker's traffic. HTTP is more common than HTTPS in a normal network and attackers would rather use the encrypted option for their traffic. It's perhaps the first time this behavior has actually been observed in a large-scale analysis of realworld networks.

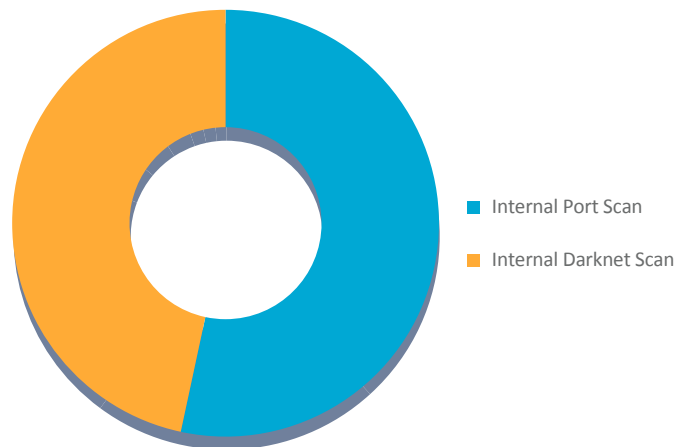
Focus on reconnaissance

Attackers spend a lot of time preparing for an attack – looking for vulnerabilities, researching employees and enumerating Internet-facing services. But once they have successfully compromised the perimeter, they are often back at square one.

Their newly compromised host usually has a vantage point into the network that was previously opaque to the attacker. In order to extend the attack, the attacker needs to know where to go, and this is the point of the reconnaissance phase.

Vectra detected reconnaissance behaviors at an alarming rate, with an increase of 270% compared to previous analysis. This points to evidence that attackers are becoming more successful at penetrating existing network defenses.

Observed reconnaissance techniques were about evenly split between internal port scans and internal darknet scans. Darknet scans leverage Vectra's long term memory of the network, and alerts can be issued when a device attempts to contact IP addresses that have not been in use. This is very useful to identify NMAP-like scans, even when they are performed slowly.



Behavior	Count	Percentage
Internal Port Scan	3288	53%
Internal Darknet	2873	47%

Figure 9: Comparison of observed reconnaissance behaviors

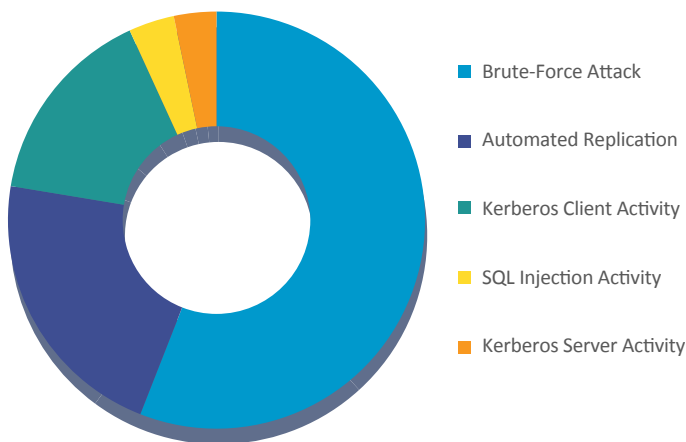
Focus on lateral movement

Lateral movement possibly represents the most strategic phase of a targeted attack. Whether through the use of spear-phishing, drive-by-downloads or plain old social engineering, attackers have proven their ability to infect individuals in almost any setting. But these initial victims rarely contain the key assets or data that the attacker is after.

Spreading from an initial beachhead to more valuable areas of the network requires lateral movement, which usually happens in one of two ways: The attacker directly attacks other clients within the network via exploits and malware or the attacker attempts to steal passwords and credentials deeper network access.

The Vectra analysis found both of these strategies in play and there was strong growth across nearly all lateral movement techniques. The most common lateral movement was also the most basic – brute force. Brute force attacks were seen across all industries but were most prevalent in the government sector.

The second most common lateral movement technique was automated replication. This detection looks for a host or hosts that deliver the same or similar exploit or payload to other hosts.



Brute-Force Attack	8721	56%
Automated Replication	3375	22%
Kerberos Client Activity	2425	16%
SQL Injection Activity	552	4%
Kerberos Server Activity	513	3%

Figure 10: Comparison of observed lateral movement behaviors

This is a very direct way to compromise additional hosts and spread through the network. Again, while this behavior was seen in most industries, healthcare networks were especially hard hit, seeing just over 50% of all automated replication detections.

A much more subtle approach to spreading within a network is to steal someone’s credentials to access additional internal resources. This can take on a variety of forms ranging from pass-the-hash techniques, to performing account scans against Kerberos infrastructure, to simply using stolen credentials to access additional services.

To identify these threats, Vectra software applies machine learning to Kerberos traffic in order to uncover anomalies. By learning the user accounts that are normally associated with specific hosts and the services they use, Vectra can determine when attackers have compromised a particular identity.

These Kerberos-based attacks were the third most common lateral movement technique and represent a 400% jump compared to the previous report. These behaviors were observed relatively equally across all industries.

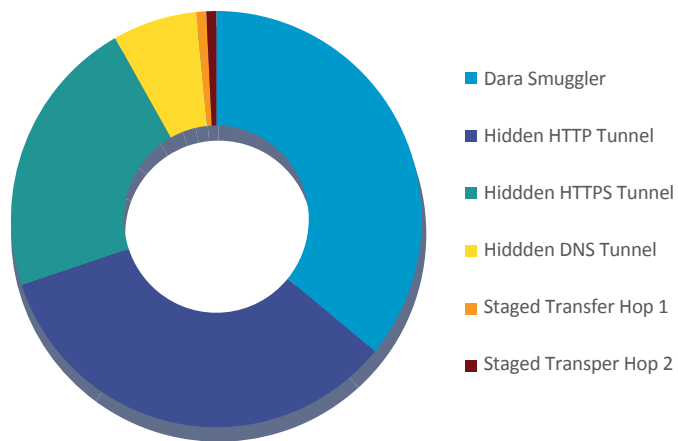
Focus on data exfiltration

While some attackers focus on destruction, the ultimate goal of most attacks is to steal data and assets. All the effort and phases of an attack ultimately lead to exfiltration. The good news is that exfiltration detections continue to be the most rarely observed phases of the attack lifecycle and organizations have been able to detect and remediate threats before a loss occurs.

The data smuggler detection was the most commonly observed exfiltration. It’s another data-driven detection that uncovers the exfiltration of information by watching and correlating the flow of data in and out of a device. This model works on encrypted channels, making it especially valuable for detecting external cyber attackers and theft from insider threats.

With exfiltration behaviors, HTTP tunnels are slightly more popular than HTTPS but still well below the normal ration of HTTP to HTTPS. The majority of these behaviors occurred within specific organizations and industries.

Media organizations showed the most data smuggler behavior, while healthcare organizations accounted for the majority of HTTP and HTTPS tunnels. DNS hidden tunnels were found predominantly in the technology sector.



Data Smuggler	455	35%
Hidden HTTP Tunnel	430	34%
Hidden HTTPS Tunnel	278	22%
Hidden DNS Tunnel	84	7%
Staged Transfer Hop 1	3	0%
Staged Transfer Hop 2	3	0%

Figure 11: Comparison of observed exfiltration behaviors

Conclusions

In this second edition of the Post-Intrusion Report, Vectra had its first opportunity to observe changes in how adversaries attack networks. The spike in reconnaissance and lateral movement within networks indicates an increased leakiness in traditional defenses.

Additionally, Vectra saw a change in attack techniques with the shift to Tor and an increased use of RATs for command and control. For the first time, Vectra was able to track the use of HTTPS as a hidden tunnel without the need to decrypt SSL/TLS, which will continue to be an area requiring further study. And for the first time, Vectra was able to measure what is common and what is unique in terms of attack patterns across industries. As sample sizes continue to grow, Vectra will continue to gain better insights into threats that networks must face.

Look for the next Post-Intrusion Report and reach out to the data science team at Vectra if you have any questions about past reports or our technologies.