NOVEMBER 2014

Post Breach Industry Report

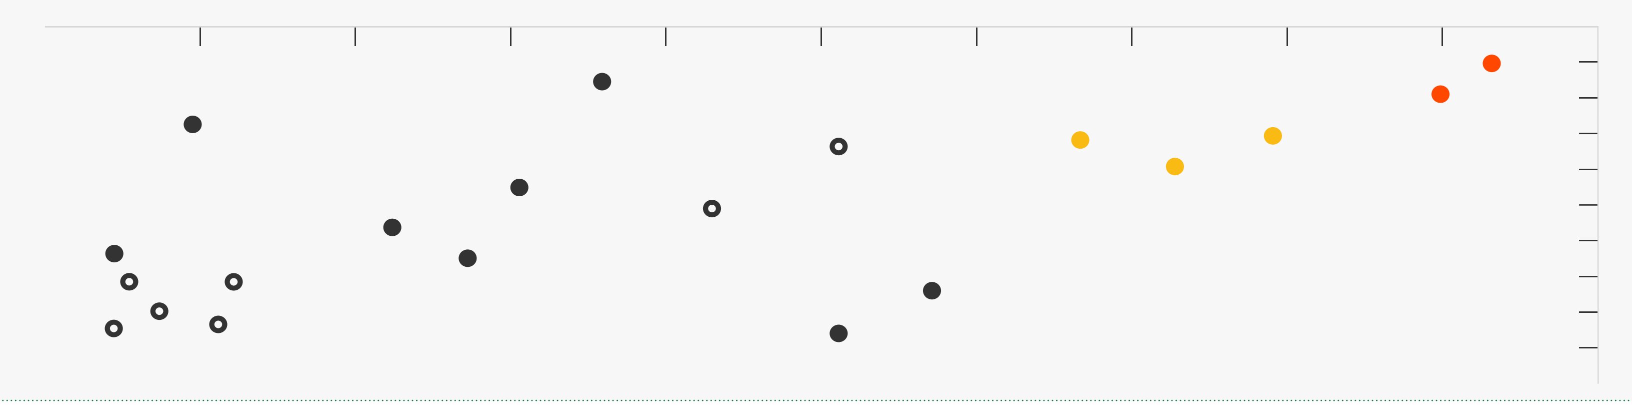

TABLE OF CONTENTS

A Real-world View of Ongoing Cyber Security Attacks

2014 is shaping up to be the year of the breach. Home Depot. Target. JP Morgan. University of California. Community Health Systems. Hardly a week goes by without bleak news of another security breach and more stolen digital valuables. Retailers, financial institutions, schools, governments, technology innovators—no organization is immune.

Cyber attacks are increasingly sophisticated and highly organized. And, they are successful despite $60 billion invested in cyber security annually worldwide. These defenses—plus the efforts of skilled information security professionals who are focused on protecting their organizations' intellectual property, honoring customer trust and upholding the law—are not working.

Anatomy of a Real-world Cyber Attack

Understanding how cyber attacks work—and what enables someone to evade strong network defenses—helps organizations better protect themselves. There's no shortage of cyber security reports that detail the volumes of threats seen by perimeter and endpoint defenses, such as firewalls, intrusion prevention systems, sandboxes and antivirus software.

But what happens after the attacker bypasses the defenses and moves into the heart of the corporate network? This first edition of the Post Breach Industry Report uses real-world data to reveal what attackers do within a network once they evade perimeter defenses.

The Post Breach Industry Report evaluates detection data from the Vectra X-series platforms deployed in production networks. Vectra Networks detects attacks at every phase of an ongoing attack, regardless of how the attack enters an organization's network and the application, operating system or device involved. The platform continuously monitors an organization's network and provides automated, intuitive and prioritized reporting so security analysts can address the highest business risks quickly. The selected organizations in this study operate in a variety of industries, including technology, financial services and higher education.

The report shows that all participating organizations had been breached by cyber attacks. Of the many tens of thousands of hosts in these organizations, more than 11,000 hosts had detections for a phase of an ongoing cyber attack. While some of these hosts were part of the same attack on an organization, 11,000 is a very high number of incidents that would have otherwise gone undetected by traditional security products.

The number of attacks detected over five months might not seem large in comparison to malware detection reports from firewall or IPS vendors. Cyber security detections on perimeter security systems occur at a higher rate than the attacks analyzed in this report, which have evaded perimeter defenses and would otherwise remain undetected inside the network.

The Vectra X-series detected multiple phases of ongoing attacks in all of the participating organizations. This report shows that 10 percent of infected computers in these organizations experienced more than one attack phase, such as command and control, reconnaissance, lateral movement and exfiltration. For hosts on which only one attack phase was detected, organizations were able to stop the attack early enough to prevent it from progressing further. The 10 percent represent hosts where the attack was stopped after the second or third attack phase detection, which still avoided or mitigated data loss.

A Crime of Opportunity or Premeditation

Most cyber attacks start with the attacker using an exploit against a vulnerability in an application or operating system. Perimeter and endpoint security products designed to detect exploits are decreasingly effective, exposing organizations to greater risk. Once an attack slips past a firewall, intrusion prevention system, sandbox or antivirus software, the attacker is free to roam around in the heart of an organization's network.

Once the attacker has gained control of a host on the network, the attack will either progress along the opportunistic or targeted paths shown below in Figure 1.

Opportunistic attacks involve many infected hosts controlled by a botnet where the attacker infects your computers to make money off of someone else. Examples include virtual currency mining, advertising click fraud and outbound denial-of-service attacks. Data from Spider.io shows that advertisers waste up to $7 million per month on fraudulent ad clicks from bots.

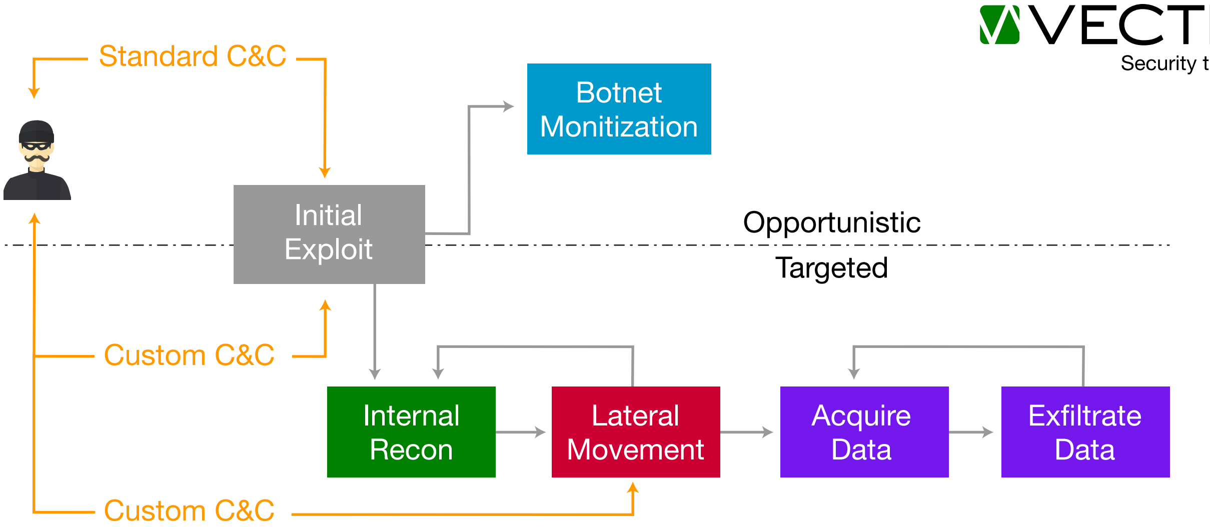

Figure 1 – The phases of opportunistic and targeted attacks.

In a targeted attack, the attacker is generally trying to steal valuable data and he will first use the first infected host to perform reconnaissance, then spread laterally to get closer to valuable data, and finally accumulate and exfiltrate the data.

The report shows that 85 percent of the hosts with multiple detections experienced an opportunisitic attack. An example is a host with detections for command and control as well as Bitcoin mining activity. Earlier this year, NSF researchers were discovered improperly using NSF-funded computers to mine Bitcoin.

Opportunistic attacks can exhibit viral characteristics. This industry report shows that 2 percent of the hosts experiencing an opportunistic attack were being used to spread botnet malware to other computers within the organization. And 5 percent of the systems showed evidence of botnet or exfiltration activity, such as moving stolen credentials off-site for use in a targeted attack. This is similar to the credentials stolen from Fazio Mechanical that were used to access a Target Corporation portal.

Targeted attacks are fewer, but are of greater concern. The data in this report shows that 15 percent of hosts with multiple detections suffered a targeted attack. Targeted attacks are rarely fully automated and almost always overseen by one or more people seeking specific information. Targeted attacks may unfold over days or weeks, and correlating multiple detections for a single host can help identify these stealthy attacks.

This report shows that 7 percent of hosts had both botnet and exfiltration detections, which indicates possible theft of credentials for use in a subsequent targeted attack against the organization. The detection of multiple attack phases resulted in 2 percent of hosts under targeted attack reaching the exfiltration stage, where the attacker was preparing to steal data and loss was mitigated due to the detection. While the number is small, the risk is real and very large.

Key findings of the first edition of the Post Breach Industry Report include:

- Cyber security detections from the Vectra X-series platform at participating organizations show that all organizations in the sample were breached
- More than 100,000 hosts were monitored within the sample and more than 11,000 hosts experienced a cyber attack
- 85 percent of attacks experienced by these organizations were opportunistic attacks
- 15 percent of the hosts in these organizations experienced a targeted attack
- All participating organizations experienced at least one targeted attack
- Of the hosts experiencing an attack, 10 percent had detections for two or more attack phases such as botnet, command and control, reconnaissance, lateral movement and exfiltration
- Even for attacks that reached the exfiltration phase, organizations had two or more opportunities to stop the attack prior to significant data loss

An In-depth Look at Detecting a Targeted Attack

One of the benefits of correlating detections by host is that these detections tell a story about the type of attack, its progression and what the attacker is doing. Figure 2 shows the series of detections for a host in a production network and an explanation of the story its detections tell.

Category	Detection Type	Host Address	First Timestamp	Last Timestamp	Threat	Certainty
Reconnaissance	Internal Darknet Scan	10.1.1.183	8/12/14 16:11	8/12/14 16:12	70	62
Lateral Movement	Brute-force Attack	10.1.1.183	8/12/14 21:18	8/12/14 13:46	65	95
Reconnaissance	Internal Darknet Scan	10.1.1.183	8/16/14 17:17	8/16/14 17:19	70	64
Reconnaissance	Internal Darknet Scan	10.1.1.183	8/16/14 22:49	8/16/14 22:56	68	61
Reconnaissance	Internal Port Scan	10.1.1.183	8/16/14 22:50	8/16/14 22:53	50	59
Command & Control	External Remote Access	10.1.1.183	8/30/14 00:27	8/30/14 00:36	82	10
Reconnaissance	Internal Port Scan	10.1.1.183	8/30/14 01:10	8/30/14 01:16	50	90
Exfiltration	Hidden Tunnel	10.1.1.183	8/30/14 19:02	8/30/14 19:04	72	95

Figure 2 – Host Detection Report for a host experiencing a targeted attack

The first observation about this host report is the time of the attack phases detected; all the detections occurred in the late afternoon or early morning hours. From this we can deduce that the host was either stationary and was being used by attackers after normal working hours or the host was mobile, but the owner of the host was traveling in a distant time zone.

The initial Reconnaissance detection is an indication of a targeted attack. The behavior detected was of an internal host that had contacted a large number of internal IP addresses that had not been recently active or were possibly never assigned. The Reconnaissance behavior is called an Internal Darknet Scan. Darknet Scans occur over longer periods than Internal Port Scans and the Vectra detection algorithm ignores contact to systems that do not respond to this host, or on this port, but which are otherwise active.

In this case, the Reconnaissance detection indicates the search was performed to build a map of the network and the connected end-user hosts and servers in an effort to learn the lay of the land. The threat and certainty scores for this Reconnaissance detection were high because of the spread of IP addresses contacted.

After the infected host built a partial map of the network, it started to exert itself laterally, increasing the attack surface, by performing a Brute-force Attack to gain access to better credentials to use to login to other systems. The infected host made many login attempts on an internal system, and the behavior is consistent with a brute-force password attack. Brute-force password attacks can be performed via different protocols (e.g. RDP, VNC, SSH, FTP, HTTP/S, SMB, SSL/TLS) and may also be a Heartbleed attack. The number of attempts and timing with which the attack was performed drives the threat score. The certainty score was driven by the total number of sessions in the attack.

The infected host was then detected performing more Internal Darknet Scans, but then the behavior switched to an Internal Port Scan, indicating that the attacker was speeding up his search. At this time, the infected host also exhibited behavior consistent with an external remote access Command and Control communication, indicating a human was controlling the attack and reinforcing that this was a targeted attack. Direct human control is never in play during an opportunistic botnet attack.

The behavior of the infected host soon changed from Reconnaissance to Exfiltration, indicating that the attacker had found data he wanted to steal, had gotten access to the data, and was beginning to send it to an offsite drop.

The Exfiltration detection was a Hidden Tunnel – a slow exfiltration method in which the infected host was communicating with an external IP address using DNS or HTTP and another protocol was running over the top of these sessions. The attack was halted within 2 minutes of the exfiltration detection, averting significant data loss.

Frequency of Opportunistic vs. Targeted Attacks

The infographic in Figure 3 illustrates the ten most common combinations of detections for hosts with multiple detections. These top ten combinations are comprised of 1,123 unique hosts, or 97 percent of all hosts that experienced two or more attack phases.

Of these 1,123 hosts, 949 (85 percent) experienced an opportunistic attack represented by the combination of both Command and Control and Botnet activity detection. With opportunistic attacks occurring on more than 80 percent of the hosts with multiple detections, detecting targeted attacks amongst the substantial noise created by opportunistic attacks is difficult. Correlating detections to the hosts under attack clears away this noise to make it easier to pinpoint targeted attacks and improve the speed of an incident response team.

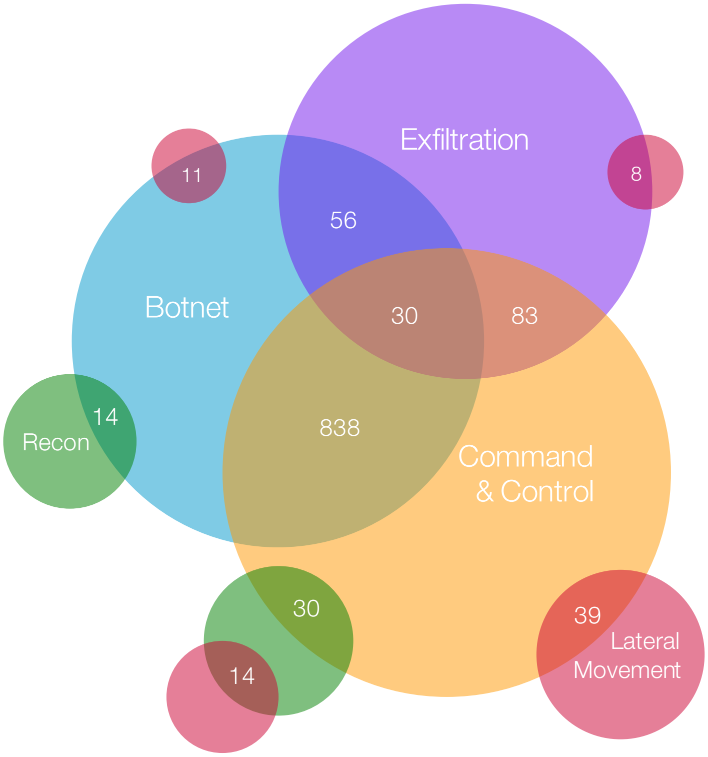

Figure 3 – Most common groupings of attack categories experienced by hosts

Although opportunistic in nature, botnet attacks that are left untended can increase in risk. Infected hosts can be instructed to spread malware, to perform Reconnaissance that seeks out system vulnerabilities, or to move laterally across machines. Of the hosts with Botnet detections, 25 hosts (2 percent) also had detections for either Lateral Movement or Reconnaissance behavior, indicating an attempt to infect other hosts.

Botnet attacks can also lead to a targeted attack. Eight-six hosts (7 percent) with Botnet detections also had detections for Exfiltration behavior where the attacker may be stealing credentials gained using either a keylogger or snooping cookies and exfiltrating them for use in a targeted attack against this organization or one of its business partners. Attackers breached Target Corp. in this manner using stolen credentials.

It was observed that an attacker may establish a hidden tunnel, which is oftentimes used for Exfiltration, as a Command and Control channel to gain greater control over a small network of machines rather than to steal data. There are 30 hosts (2.6 percent) with a combination of Command and Control, Botnet and Hidden Tunnel Exfiltration detections that fall into this category. Hosts that lie at this intersection may represent a more significant risk, even if the patterns of detections do not suggest theft of a large volume of data.

Lateral Movement, Reconnaissance and Exfiltration are indicators of targeted attacks and approximately a tenth of the hosts had a combination of Command and Control and detection of Reconnaissance, Lateral Movement and/or Exfiltration behavior.

Command and Control – Human Factor Tips Off Targeted Attacks

Having greater insight into Command and Control detections enables organizations to distinguish between human-driven targeted attacks and automated opportunistic attacks. Targeted attacks are rarely fully automated, rarely use a command-and-control server that has been used before, and are often carried out by individuals seeking specific information.

While identifying the type of Command and Control detection alone cannot be used to make definitive claims about the nature of an attack, External Remote Access and The Onion Router (TOR) detections are the most common detections of a human-driven targeted attack.

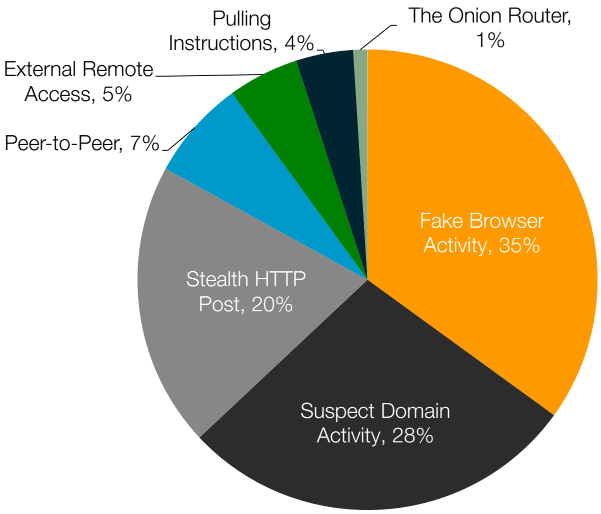

Figure 4 – Frequency of different types of Command and Control detections

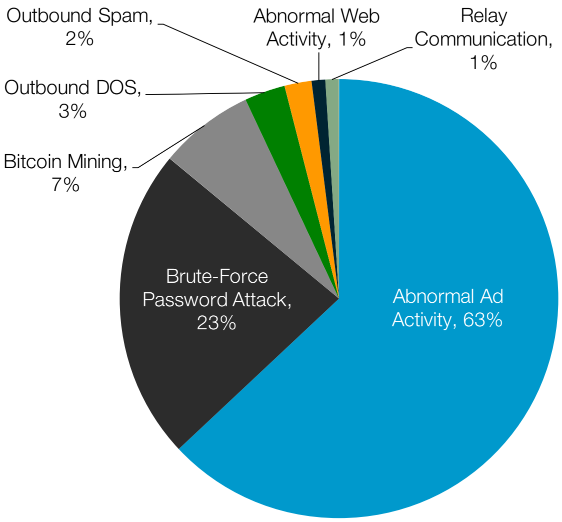

Figure 5 – Frequency of different types of Botnet detections

External Remote Access detections represent 5 percent and TOR detections represent 1 percent of all Command and Control detections from the sample. These 6 percent of detections further demonstrate that targeted attacks are less frequent than opportunistic attacks and can be hidden by the noise caused by opportunistic attacks.

How Opportunistic Attackers Make Money and Harm Your Reputation

Botnet activity is a clear indicator of an opportunistic attack in which the bot herder is utilizing the infected host computer, its network connection and, most of all, the unsullied reputation of the organization's IP address to make money. Opportunistic attacks present several risks to an organization. All opportunistic attacks create noise that may disguise higher-risk targeted attacks and they can also cause the organization's IP address to end up on a blacklist. The compromised host may also be instructed to perform a direct attack on the organization.

The variety of possible Botnet behaviors should be considered for the risk they represent to the organization. Abnormal Ad Activity is usually detected when a user is tricked into installing malware that monetizes ad clicks. Brute Force Password Attacks on other organizations are detected when malware looks for other Internet-accessible systems to breach in order to extend the botnet's footprint. In this scenario, an organization's computers are being used to actively probe for vulnerabilities in another organization's perimeter defenses.

When Bitcoin mining is detected, it is important to understand if the user has intentionally installed cyber currency mining software. The risk caused by virtual currency mining may be minimal, though such a user may also be prone to installing other money-making software, which may introduce a higher risk.

Detections of Outbound DoS, Outbound Scans and Outbound Spam may materially affect the bandwidth available for legitimate functions, reducing productivity.

Relay Communication was detected in at least one of the deployments and these hosts were being used in a distributed denial of service attack on another organization.

Organizations should consider that users who become infected with botnet malware represent a more general risk since they may also become the entry point for more harmful malware.

Reconnaissance – How Targeted Attackers Build a Map of the Network

When a vulnerability is exploited and malware is first installed on a host, the attackers rarely land directly on the host that holds the data they are trying to steal. The initial malware installed is like a person who has been beamed into a dark building. The attacker uses the infected host to perform reconnaissance, similar to a person in a dark room feeling around for walls, unlocked doors and corridors to find his way to an exit.

To perform reconnaissance, the attacker uses the infected host to create a map of the network and the connected end-user hosts and servers so he can determine the location of the data he wants to steal.

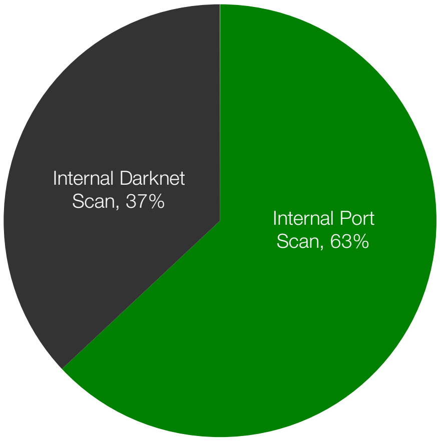

Figure 6 – The frequency of Darknet Scan vs. Port Scan detections for reconnaissance

Reconnaissance can take two forms: slow and stealthy or fast and broad.

Darknet Scan detections indicate than an internal host has slowly contacted a large number of internal IP addresses that have not been active in any way in the recent past. Internal Port Scan detections occur when an internal host has either quickly attempted contact with a large number of internal IP addresses on a small number of ports – a network scan – or with many ports on a small number of internal IP addresses – a host scan. Given that they target slow and stealthy reconnaissance, Darknet detections take a greater amount of time to trigger than internal port scans.

Ten percent of the hosts with Reconnaissance detections in this report included a combination of Internal Darknet and Internal Port Scans. On these hosts, the attacker likely performed a slow-and-stealthy Darknet Scan to first create a macro-level map of the network and hosts, then the attacker switched to a fast-and-broad Port Scan to perform more detailed reconnaissance on specific hosts.

How Targeted Attackers Spread Laterally to Get Closer to High-value Data

Lateral Movement within a network exposes a larger surface to the attackers and exposes organizations to substantial risk of data acquisition and Exfiltration. An attack spreading laterally can involve attempts to steal account credentials or to steal data from another machine. It can also involve compromising another machine to make the attacker's foothold more durable or to get closer to target data.

The majority of Lateral Movement detections in this report are Brute-Force Password Attacks. This detection type can be triggered when the infected host makes many login attempts on another internal system. Note that the successful harvesting of login credentials – usernames and passwords – of accounts, particularly more privileged accounts, is a typical stage in the progression of a targeted attack. Even false positives due to the misconfiguration of authorized applications can create significant stress on internal systems and were usually corrected by IT staffs.

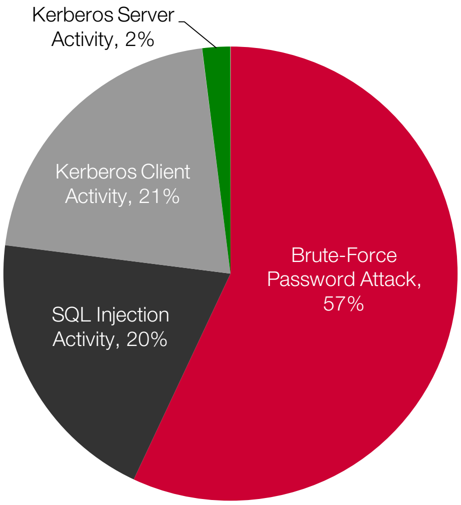

Figure 7 – Frequency of Lateral Movement detections observed

The second most frequent detection at 20 percent, SQL injection, is triggered when an internal host sends a series of requests to a Web server and embeds SQL commands into HTTP Post or URL data to gain access to a company's backend database. Probing and potentially exploiting an internal web application's vulnerabilities can be a prelude to the stages of a targeted attack in which the attackers gain access to and then exfiltrate data. Application software that relies on passing fragments of SQL statements over the network may well be vulnerable to attackers who can feed the applications input that varies greatly from what the applications are designed to handle.

Kerberos Client and Server Activity attacks represent nearly a quarter of all Lateral Movement detections and together form two sides of the same coin. Kerberos Client Activity is detected when a Kerberos client (e.g., an application acting on behalf of an end user) attempts a suspicious amount of authentication or service requests using either a small number of services and accounts (brute force) or a large number of services and accounts (scan). A Kerberos Server detection arises when the server denies a suspicious amount of authentication requests from multiple clients using multiple services. Misconfigurations that cause the equivalent of brute-force attacks should be fixed by IT staff as they could adversely affect the performance of both the client and the server, and hide real attacks.

Summary

All cyber attacks start with an initial exploit. As this report shows, perimeter and endpoint security designed to detect exploits are decreasingly effective at stopping this initial incursion, exposing organization to greater risk. Once an attack slips past a firewall, intrusion prevention device, sandbox or antivirus software, the attacker is free to continue unabated in the heart of an organization's network.

Regardless of whether the initial exploit can be detected, the subsequent actions of an attacker are very consistent. The attackers' goals may be opportunistic where the vulnerability provides an opportunity for him to dramatically expand the footprint of a botnet to increase his income. Or, the attackers' goals may be to attack your organization and steal high-value data.

Detecting the attackers actions once inside the network provides multiple opportunities for organizations to stop both opportunistic and targeted cyber attacks. Correlating detections to the host under attack will tell a story of what the attacker is doing and help triage targeted attacks from opportunistic ones, increasing the effectiveness of an incident response team.

All the organizations participating in this report experienced both opportunistic and targeted attacks. Opportunistic attacks outnumbered targeted attacks by a factor of 3 to 1. Ninety percent of these attacks were stopped after a single attack phase was detected. And, 10 percent of the hosts attacked in this study experienced two of more attack phase detections prior to the attack being stopped.

All of the attack phases detected in this report are ones that evaded organization's perimeter and endpoint security systems. The detections by Vectra X-series platforms provides both visibility into the attack and multiple opportunities to stop it or mitigate data loss if the attack progresses to the exfiltration phase. Even for attacks that reached the exfiltration phase, organizations had one or more opportunities to stop the attack prior to significant data loss.