KENNASECURITY.COM

How the Rise in Non-Targeted Attacks Has Widened the Remediation Gap

September 2015

Breaches are on the rise, and every company – from SMBs to enterprises – is at risk. The reason for this intensifying inflection is in part non-targeted attacks, which come at such scale that security teams can't keep up with them.

Non-targeted attacks represent a vastly different challenge than the more widely publicized Advanced Persistent Threats: attackers in volume care less about who they hit and rather what they can get, which is why they also deploy automated methods that give them economies of scale. They can go farther, and hit more – all in hopes of finding data they can use (credit cards, SSNs, etc).

When attackers employ a non-targeted attack, they are looking for specific vulnerabilities that they have the ability to exploit, as opposed to specific companies to breach. It's like going to as many homes as possible and rattling windows in order to find one that they can pry open. If they're successful – and they find something that they can exploit – then they will extract as much information as they can. For many organizations, non-targeted attacks are much more dangerous than the highly publicized breaches that the press tends to report.

The much-publicized Heartbleed exploit is a good example of a non-targeted attack due to the ease with which an attacker could exploit multiple targets at once. Unlike a crafted, targeted attack against a specific enterprise, attackers scanned the Internet-vulnerable websites, and engaged in "spray and pray" in hopes of being able to extract juicy bits of information. Of the organizations that Kenna has evaluated, 100 percent are susceptible to vulnerabilities – which correlate to at least one stable publicly available exploit.

The sheer increase in volume of automated attacks over the last year has created a "new normal" where security teams can't keep up with the signal-to-noise ratio. James Trainor, acting assistant director of the FBI's Cyber Division, claims that data breaches are up 400 percent in 2015, and the workforce for the cyber division needs to be doubled or tripled.[1] A PWC report saw an almost 100 percent increase year over year in 2014.[2]

The research of Kenna Security attributes this increase not to the sophistication of attacks themselves, but rather the sophistication of the attackers' modus operandi: they are getting better at automating their attacks. The result is an unprecedented volume of attacks as well as volume of businesses exposed to these attacks. Due to the inability of Information Security teams to match the pace of automated attacks, a significant gap has appeared in the time that critical vulnerabilities appear and the time it takes for security teams to fix those vulnerabilities.

In this report, we will examine:

- **How the rise of non-targeted attacks have contributed to the remediation gap**
- **An example of serious vulnerabilities that have gone unremediated**
- **Recommendations for closing the remediation gap**

[1] http://thehill.com/policy/cybersecurity/242110-fbi-official-data-breaches-increasing-substantially

[2] http://www.pwc.com/en_US/us/private-company-services/publications/assets/pwc-gyb-cybersecurity.pdf

Methodology

This report does not examine attacks from the entire population of known companies, but rather a sample size of 50,000 organizations, 250 million vulnerabilities and over one billion breach events – with data taken from 2014-September 2015. The data from these organizations is derived from Kenna's threat intelligence data as well as data from various partners, such as AlienVault, Dell SecureWorks, Verisign, SANS ISC and US-CERT.

When analyzing targeted attacks, we collected real-time data by looking for successful attacks. We then performed a correlative analysis of exploits being fired off, and examined indicators of compromise. Finally, we pulled data from partners with forensic teams.

In regards to analyzing non-targeted attacks, we used seven data feeds through our platform with access to 10 million successful attacks/week – most notably proprietary data collected through AlienVault's Open Threat Exchange (OTX). By executing this approach, we were able to estimate the probability that a vulnerability might be exploited, as well as the sheer volume of attacks, based on the volume of attacks displayed by the aggregated data.

The data we collected provides strong insights and visibility into the remediation gap found in 2014 through September 2015, and why this gap goes under-reported. The highly publicized breaches that the press tends to report—i.e., the Ashley Madisons of the world that tend to grab the headlines.

Key Findings

- **On average, it takes businesses 100-120 days to remediate vulnerabilities.**
- **At 40-60 days, the probability of a vulnerability being exploited reaches over 90 percent – indicating that most successfully exploited vulnerabilities are likely to be exploited in the first 60 days. The gap between being likely exploited and closing a vulnerability is around 60 days.**
- **In 2015, as of August 1, 2015 there have been a total of 1,272,152,215 successful exploits from our sample size of approximately 50,000 organizations. This is compared to 219,951,631 exploits in 2013 and 2014 combined.**

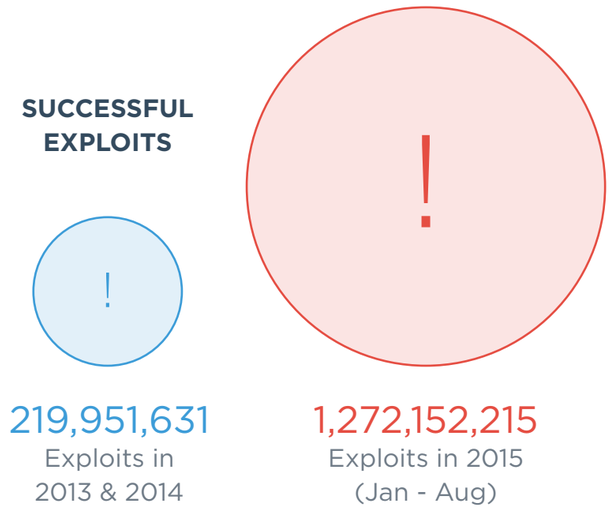

Companies Leave Vulnerabilities Unpatched for Too Long

In the first year when a vulnerability is released, it's likely to exploited within 40-60 days. However, it takes security teams between 100-120 days on average to remediate existing vulnerabilities. Therefore, the gap time between likely being exploited and closing a vulnerability is 60 days. Security teams are (1) slower than the attackers, (2) they often remediate vulnerabilities which aren't being exploited, and (3) they're being overrun by non-targeted, automated attacks.

When a vulnerability first appears, attackers move extraordinarily fast and learn quickly how best to exploit it. Over time, they continue to get better at exploiting the vulnerability – and around the 40-60 day mark, the probability that a vulnerability will be exploited hits an inflection point and is around 90 percent. This means the majority of these breaches occur in the first 40-60 days.

Cumulative Probability of Exploitation

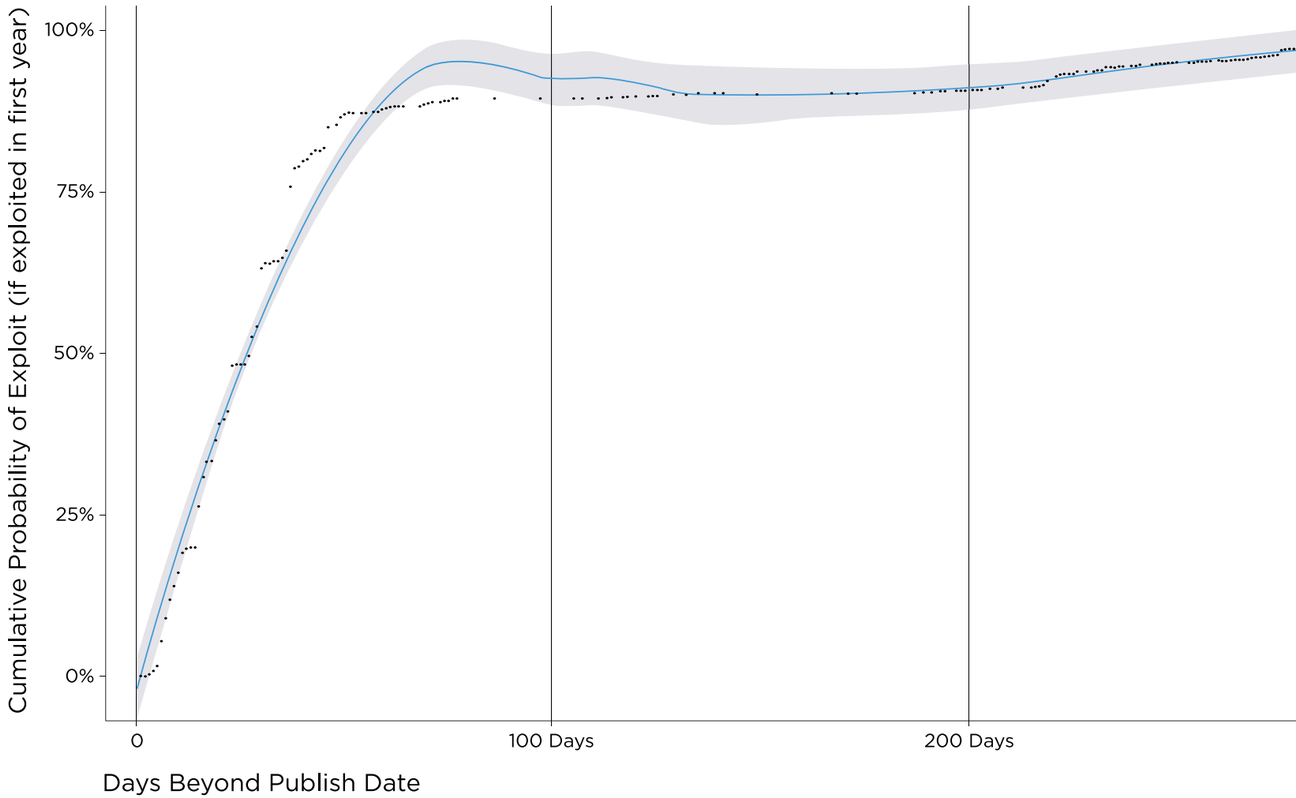

Figure 1: Cumulative Probability of Exploitation – This shows the probability that a CVE that is exploited in the first year will be hit X days after publication. At 40-60 days, that probability is over 90 percent.

Security teams can't move at the same speed as their adversaries. The graph below demonstrates the average monthly days that it takes to close a vulnerability once it's been discovered. Keep in mind that this is only the vulnerabilities that have been closed, meaning that there's a large body of vulnerabilities that haven't been remediated – which drives the real averages even higher.

Average Days to Close by Month Quadratic Fit

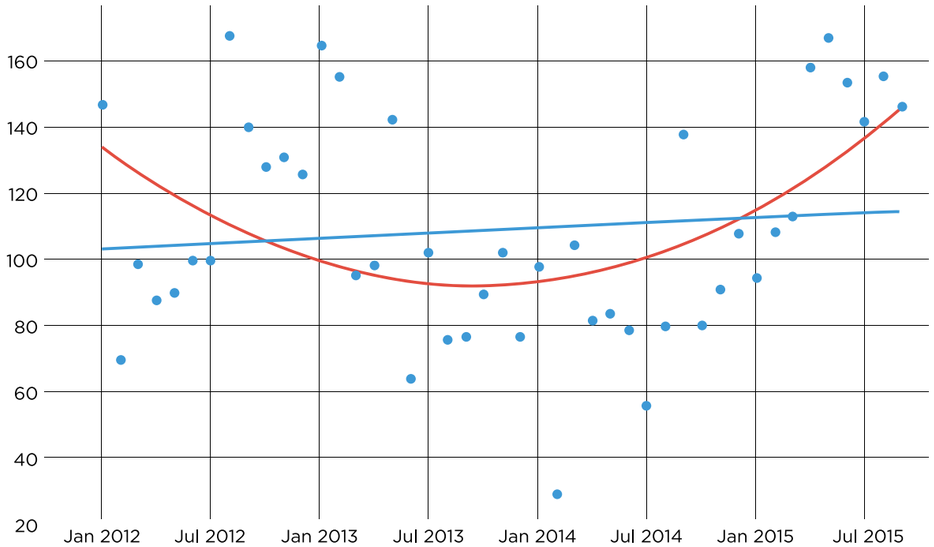

Figure 2: Quadratic and Linear fits of by month average days to close a vulnerability across Kenna Security's 2000 customers (sample size = 1,668,000 vulnerabilities, period = 2012-2015). Both fits indicate mild growth over time (not statistically significant), but stable clustering around 100-120 day.

Therefore, a high probability exists that every organization – from SMBs to enterprises will have a high degree of critical, unremediated vulnerabilities. You don't have to have incurred the ire of a rogue government, or group of hackers out for blood, in order to be a victim; you simply have to have the requisite (and highly common) vulnerabilities that match at least one non-targeted attack.

The gap time (on average) between detection and close is 60 days. This represents an estimate of the half-life, not of the time to close. Based on our analysis, the length of time a company has to react to vulnerabilities before attackers strike is 40-60 days after the release for well-known vulnerabilities. At the 40-60-day mark, the probability that an organization has already been hit by a non-targeted attack is significant.

The 40-60 day number is an estimate, but nonetheless represents a very real danger. The bottom line is that vulnerabilities open to non-targeted attacks are, by and large, not being remediated quickly enough to prevent widespread exploitation.

> *At the 40-60-day mark, the probability that an organization has already been hit by a non-targeted attack is significant.*

3 Examples of Vulnerabilities Commonly Exploited by Non-Targeted Attacks

When we talk about unremediated vulnerabilities that fall prey to attacks at scale, one of the points we need to make is that the vulnerabilities in question are often very old, well-known weaknesses that simply haven't been fixed yet. We've seen this over and over again as we evaluate the data.

In many cases these vulnerabilities are not sexy, and they don't hog the spotlight – but in many environments they actually represent major weaknesses. Below we've provided three examples of these vulnerabilities simply as a point of consideration. According to our research, these three vulnerabilities have each been exploited over 100,000 times in 2014 alone. The scale alone is an indicator that the attack is automated, launched at scale.

The vulnerabilities we want to highlight are CVE-2010-3055, CVE-2002-0649, and CVE-2000-1209. Let's name them, as though they were as publicized and talked about as a vulnerability such as Heartbleed.

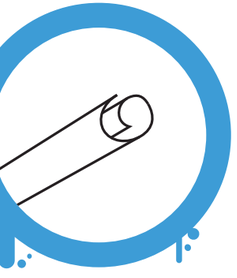

1. CVE-2010-3055 was exploited 121,000 times in 2014. Let's call it the **Poster** vulnerability. It allows attackers to run arbitrary code in phpmyadmin via a POST request, and phpmyadmin runs millions of sites worldwide. It's a CVSS 7.5, which means it's bound to fly under the radar more often than not. But it shouldn't. Security teams need to start worrying about Poster.

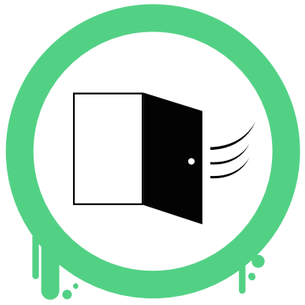

2. Let's call CVE-2002-0649 the **Slammer** vulnerability. It's an ancient worm that exploits SQL Server 2000 and Microsoft Desktop Engine 2000. Reading the Wikipedia article on the worm makes it seem like it's a long forgotten problem, but we witnessed 156,000 successful exploitations in 2014. It's not new, it's not hip, it's not current, so one talks about it – but it's a significant threat.

3. Last up is **Enterprise**, which exploits (1) Microsoft SQL Server 2000, (2) SQL Server 7.0, and (3) Data Engine (MSDE) 1.0, including third party packages that use these products such as (4) Tumbleweed Secure Mail (MMS) (5) Compaq Insight Manager, and (6) Visio 2000, and are exploited by the Voyager Alpha worm. CVE-2000-1209 is also not to be forgotten, with 272,000 successful exploitations. Resistance is futile?

Here's more info about each vulnerability:

https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2010-3055
https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2002-0649
https://web.nvd.nist.gov/view/vuln/detail?vulnId=2000-1209

These vulnerabilities are not new – in fact, they're extremely old – and yet they perfectly represent the kind of unremediated vulnerabilities that automated attacks attempt to find. They're the windows that the criminals rattle around and try to pry open.

How long it takes for attackers to exploit a known CVE, and most importantly, for a company to be hit by an exploit depends on the type of vulnerability. The probability of successful exploitation depends on the type of vulnerability – for firewall, server or workstation vulnerabilities the timelines differ significantly.

They also differ for the highly publicized vulnerabilities, but here the timelines are fairly quick – the first few see significant activity, other than Heartbleed, which took a while to detect and a while for the successful exploitations to start occurring.

Recommendations for Reducing the Remediation Gap

What can organizations do when the number of vulnerabilities keep growing, and they don't have enough resources to close them before that dangerous 40-60 day mark?

For one thing, it's not enough to throw people at the issue anymore. And it's not like most companies have the headcount budget to do that anyway, or access to a deep enough InfoSec talent pool. To properly combat today's non-targeted threats, security teams need to start using the same kinds of automated methods that their adversaries are using. They need to perform their tasks using computational models and algorithms to gain the scale they need to find and remediate critical vulnerabilities.

Clearly, peering into CVSS scores and crunching Excel data won't get the job done quickly enough to avoid that 40-60 day delay – and, of course, in many cases a successful breach might occur much more quickly than 40-60 days.

It's necessary to give security teams an unfair advantage in determining what vulnerabilities to remediate. The best resources will use up-to-the-minute threat intelligence to determine exactly what is being exploited at the current time, and which ones specifically pose a threat to a particular organization. There's no need to close 50,000 vulnerabilities – just the relatively small number of vulnerabilities that actually pose a real and present danger. Doing so must be done quickly, and at scale.

In addition, it's important not to be distracted by bright and shiny objects. The vuln-du-jour – be it Heartbleed, Poodle, Venom, or some unnamed media darling code name that we'll all find out about soon – may not be the real threat (although those vulnerabilities continue to cause problems, as our "Predictions" sidebar shows). The real threat may be some very old, unsexy vulnerability that's ready to be exploited by a non-targeted attack. It's that relatively unremarked upon vulnerability that should be remediated as quickly as possible.

Companies that do not adequately prepare for non-targeted attacks are most at risk, and therefore need to increase their awareness and implement solutions that monitor exposure to breaches through active vulnerability analysis. One way to do this is to stop beefing up the headcount budget, and instead use automated methods that enable the rapid prioritization and remediation of critical vulnerabilities.

Attackers who are looking for information at scale have long abandoned manual methods. Unless InfoSec teams do the same, they won't be adequately prepared to protect their organizations.

Predicted uptrend for Heartbleed
We predict around 5000 successful exploitations of Heartbleed per day for the next month from the publication date of this report (September 22, 2015). Link to the forecast and six-month historical data here - **http://www.stathat.com/s/OC04Jh79kPpR**

Predicted up-trend for CVE-2005-1256
We predict an uptrend a remote authenticated code execution vulnerability in IMAP daemon imapd32.exe, and predict around 2000-5000 successful exploitations per day for the next month. - **http://stathat.com/s/8AMpoogUsDQM**

Predicted flat trend for Poodle vulnerability
We predict a flat trend around 90,000 successful exploitations per day for the Poodle vulnerability for the next month, based on weekly data with strong recurring weekly counts - **http://stathat.com/s/WDPPwRfYIruM**

About Kenna

Kenna is a Risk & Vulnerability Intelligence platform that correlates external Internet breach data, exploit data and zero-day threat intelligence with internal vulnerability scan data so organizations can focus on fixing the most critical vulnerabilities. Kenna processes over a billion vulnerabilities a day against Internet breach data for its users.

For more information, visit kennasecurity.com.