REPORT

2016 State of Vulnerability Risk Management

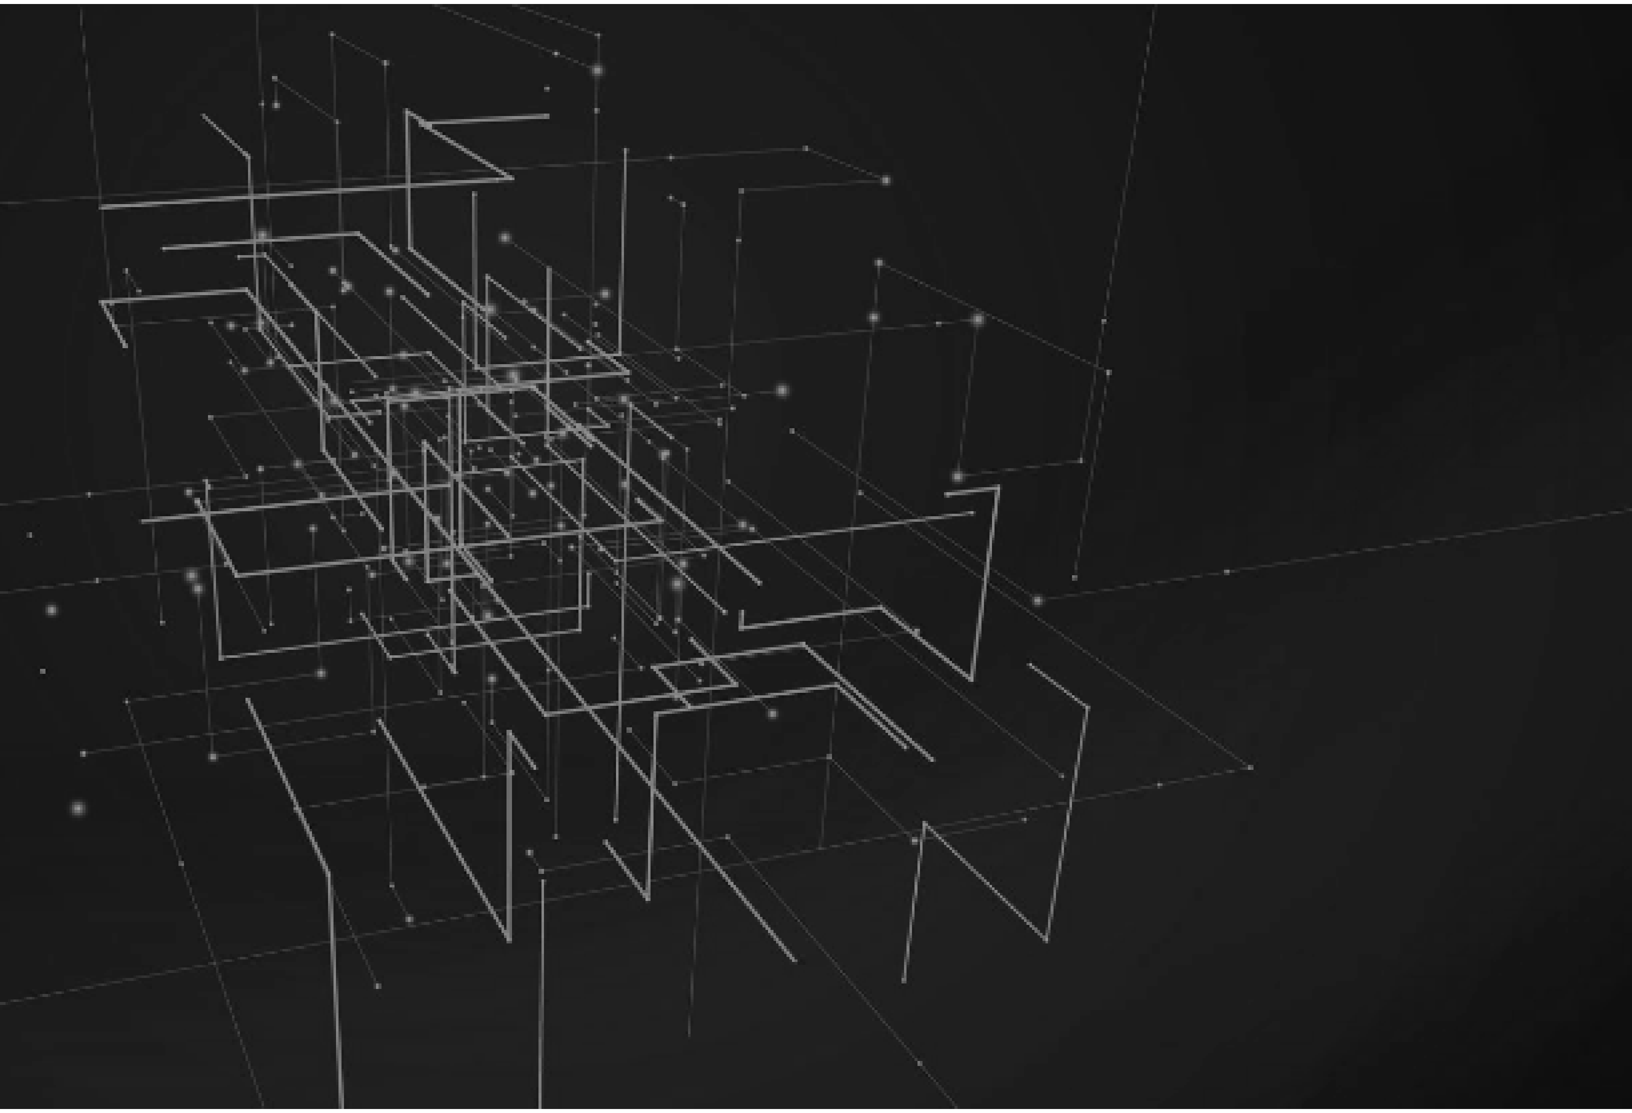

NOPSEC

nopsec.com | hello@nopsec.com

Table of Contents

Introduction

Certainly, the vulnerability management situation in 2016 has evolved and shown some signs of improvement. However, for far too many organizations, the approach to vulnerability management is stuck in the past. Scanning is aimed only at reporting for compliance and remediation of security vulnerabilities is sporadic at best. Risk management of assets and related vulnerabilities are frequently unsystematic or merely use the Common Vulnerability Scoring System (CVSS) score, leading to an incomplete measure of vulnerability risk.

This predicament is the outcome of multiple, intertwined factors. By and large, vulnerability management programs are driven by the need to document and report on vulnerabilities for compliance purposes, and measured in terms of lower overall vulnerability count. In turn, the current methodology for classifying vulnerabilities based on criticality lacks the necessary context to establish a set of actionable priorities and lacks predictive value. Scanner output is so vast that it leads to information overload and prevents organizations from quickly moving from vulnerability detection to remediation. As vulnerability management programs remain time-consuming and manually intensive for many organizations, they miss crucial steps in remediation: applying patches and making security configuration changes.

Measuring vulnerability risk is a much more intricate and nuanced activity than simply considering the vulnerability’s CVSS score.

From the outset, NopSec has focused on pioneering a way to measure vulnerability risk based on threats to the organization’s valuable assets in a hypothetical event of a breach. As presented in this 2016 State of Vulnerability Risk Management Report, measuring vulnerability risk is a much more intricate and nuanced activity than simply considering the vulnerability's CVSS score.

The Information Overload Antidote – Prioritization Through Realistic Risk Scoring

Our research indicates that building a sustainable and repeatable approach to prioritization of vulnerabilities requires more than evaluating the vulnerability's Common Vulnerability Scoring System (CVSS) score in isolation — especially if the output generates a high volume of 'critical' vulnerabilities. Weighting all critical vulnerabilities with equal risk has the practical outcome of prioritizing none. Vulnerability management programs must have a prioritized set of vulnerabilities, driven by insights into the relative risk to the organization, to operate effectively. The challenge is compounded by the relative lack of visibility into asset infrastructure and frameworks to assign a business value to the asset where vulnerabilities have been identified.

The wide variance in the exploit window illustrates the importance of an automated and consistent vulnerability remediation program that stays up-to-date with the latest exploit information.

By incorporating context and additional data feeds, including social media trend analysis on exploits, organizations can move beyond information overload and advance a risk-driven approach to vulnerability remediation.

Revisiting Risk Score Assumptions and Redefining Critical Vulnerabilities

The rise of the cybercrime economy means that tools are widely available to spot vulnerabilities and exploit them without requiring an enormous amount of technical sophistication on the part of an individual attacker. Critical vulnerabilities represent low-hanging fruit for opportunistic hackers using automated tools to execute exploits, and as such should be urgent priorities for remediation. However, we at NopSec have analyzed the CVSS scoring system and concluded that it does not accurately represent critical vulnerabilities. Instead, NopSec has set out to devise a risk scoring methodology more representative of the current threat environment organizations face. It is built on a multidimensional model that integrates scanner results with external data feeds. Our Technical Risk Score re-weights CVSS attributes based on our research (weighing the factors correlated with attacks and data breaches more heavily) and incorporates additional data about public exploit availability, malware correlation, and social media feeds. Additionally, NopSec's risk score measures the "business risk" of a vulnerability by taking into account the context of the information asset a vulnerability affects.

The goal of this report is to shed some light on the current threat landscape for organizations, assess the strengths and weaknesses of current vulnerability evaluation systems such as CVSS, and explore additional metrics for determining the risk of a vulnerability.

Data and Methodology Overview

The analysis in this report is based on aggregated anonymized NopSec Unified VRM client data, which consists of over 1,000,000 unique vulnerabilities found on our clients' systems. For our purposes, we define a unique vulnerability as a unique combination of client, vulnerability ID, asset, and port affected. We use this definition because a vulnerability's intrinsic attributes are only one part of risk – the context in which a vulnerability is present is often just as important.

Our clients span a wide range of industries, but for the purposes of this report, we have classified them into one of five broad industry categories: Financial, Technology, Healthcare, Insurance, and Other. We have integrated our client data with information on public exploits (from sources such as the Exploit DB and Metasploit), malware correlation data, social media information from Twitter, and other sources (such as CWE and CPE information) to add additional context and give a comprehensive overview on the State of Vulnerability Risk Management for our clients.

It is important to note that our analysis comes from a convenience sample of our clients – as such, we do not claim that this is a definitive analysis of all possible threats. The possibility of sample bias exists, and this should be kept in mind throughout the report. However, we believe that our research offers important insight into how companies in various industries address vulnerabilities, universal weaknesses companies across industries share, and factors that should be incorporated into a comprehensive threat detection and remediation program.

Vulnerability Counts

We begin our analysis with an overview of vulnerability counts for clients in each of our industry categories. By examining overall and per scan vulnerability counts, we can gain insight into the magnitude of threats clients in different industries face, and overall remediation trends.

A statistical note: because the distribution of vulnerability counts for clients is right-tailed (there are a few companies with very large numbers of vulnerabilities), we found that a few clients had an overly large impact when we averaged our data. Therefore, we chose to take the median, as it is significantly more robust to outliers. We will continue to use the median of our data throughout this report, as we believe it gives a better representation of the threat landscape for a “typical” client.

Below is a graph depicting the median number of vulnerabilities discovered per client. These numbers illustrate the total number of vulnerabilities discovered on typical NopSec client systems since they began using Unified VRM.

Count of Total Vulnerabilities Discovered per Industry

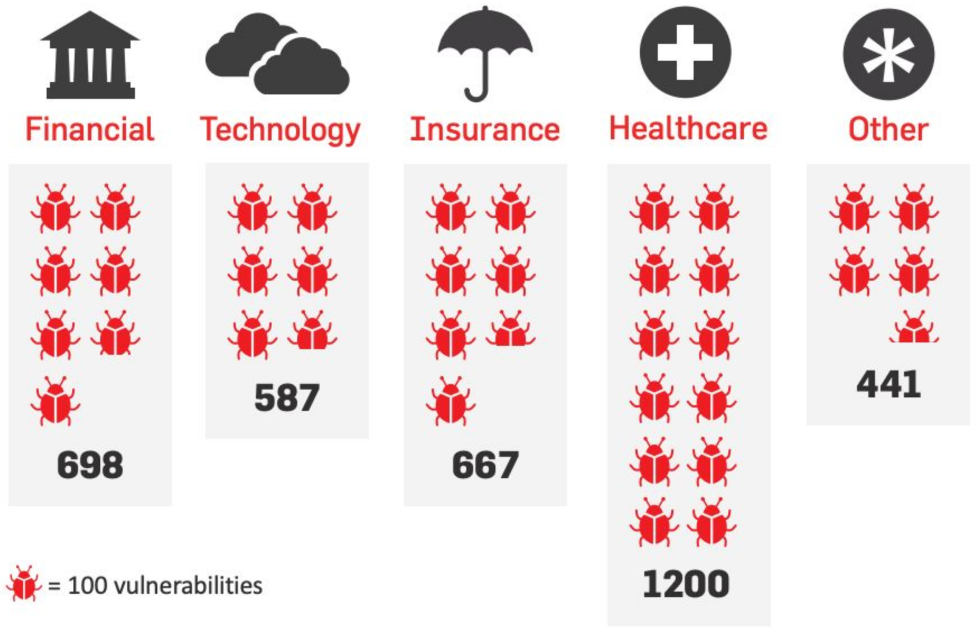

Vendor Analysis

Next, we examine top vendors by industry in order to determine which are the most vulnerable. Again, in order to get the best picture of what a “typical” client faces, we measure the median number of vulnerabilities each client has. The chart below shows the vendors with the most vulnerabilities per client associated with them.

Vendors with the Most Vulnerabilties per Client

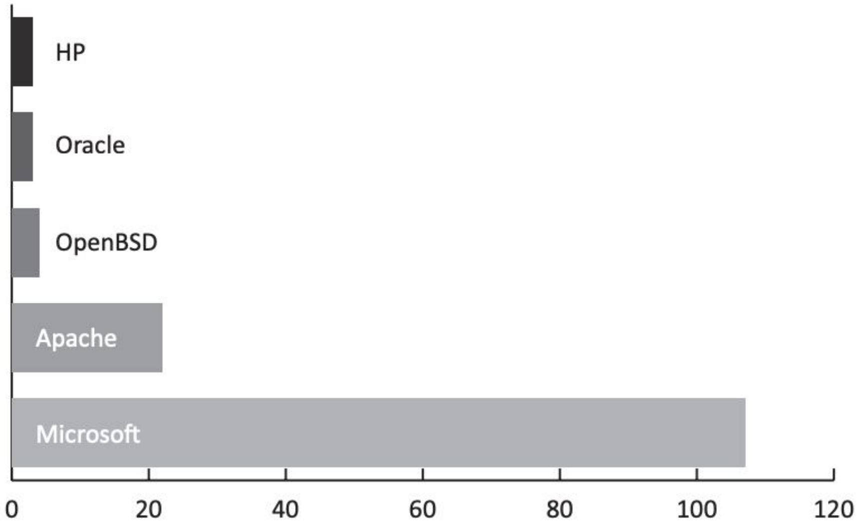

Breaking these numbers down by industry gives greater insight into the most vulnerable vendors – namely, that clients in different industries have very different top vendors.

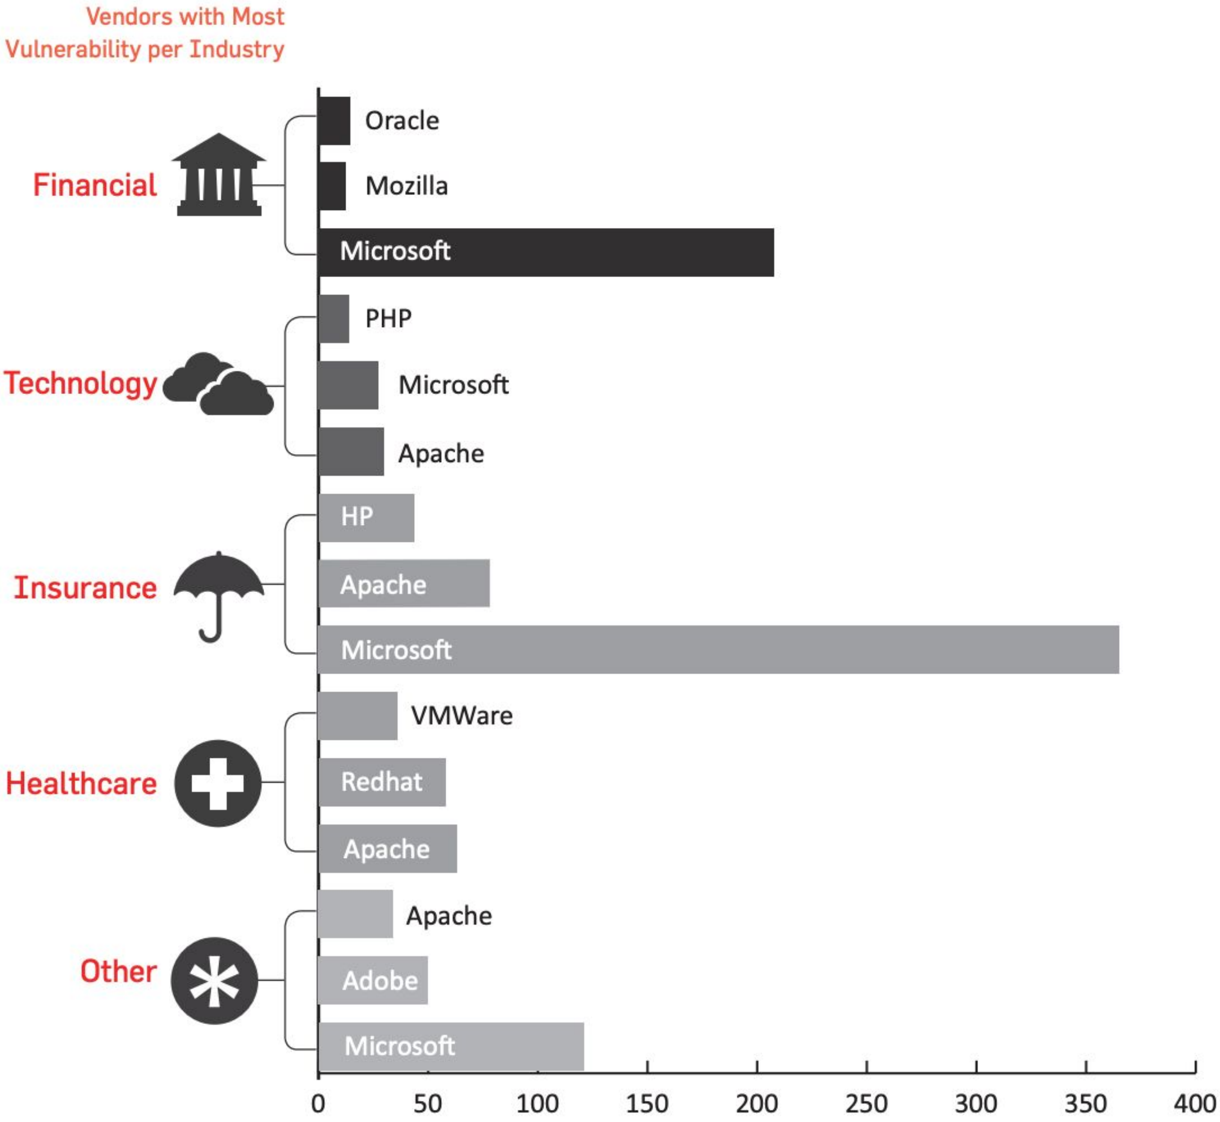

NopSec judges that these numbers are more an indicator of widely used products by industry more than an indicator of which vendors are most vulnerability-prone. Nevertheless, these numbers yield some interesting insights and practical actions for organizations looking to improve their vulnerability prioritization and remediation. While it is no surprise that Microsoft is one of the top vendors by vulnerability count in every industry vertical due to its wide global deployment, application vulnerabilities such as Adobe, Mozilla, and VMWare are also significant. This growing prominence of application vulnerabilities, especially when taking into account wide use of Microsoft software for both workstations and servers, indicates application-oriented patch management should be more consistently integrated into patch programs.

Security Vulnerability Weaknesses Analysis (CWE – Common Weakness Enumeration – Analysis)

In this section, we examine the most common weaknesses for clients across industries. When combined with the vendor analysis from above, this analysis can provide important information on unique challenges for each industry as well as commonalities across industries.

From the MITRE definition – "CWE – Common Weakness Enumeration — provides a unified, measurable set of software weaknesses that is enabling more effective discussion, description, selection, and use of software security tools … as well as better understanding and management of software weaknesses related to architecture and design."

The graph below illustrates the 5 most common weaknesses found on our clients' systems. These numbers reinforce the importance of remediating application-layer vulnerabilities. While user and service permissions, privileges, and access controls account for the highest number of vulnerabilities per client, well-known and longstanding application-layer vulnerabilities such as cross-site scripting (XSS) and other forms of improper input validation vulnerabilities represent a significant proportion of the overall weakness count.

Five Most Common Weaknesses on Client Systems

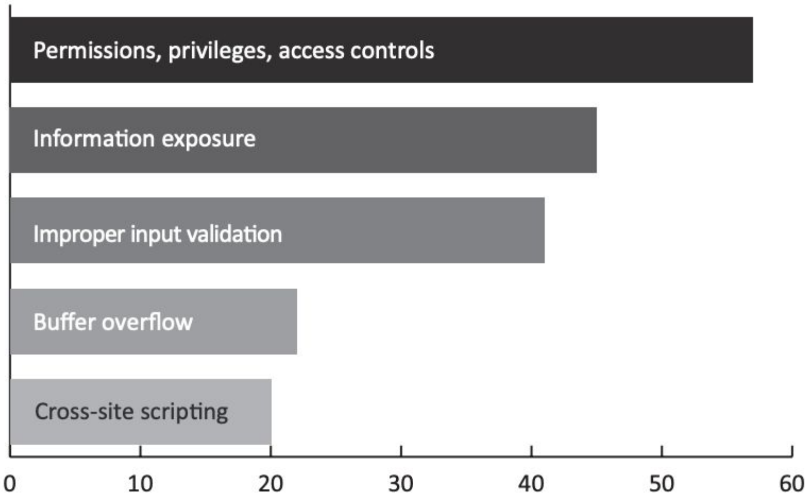

Our analysis by industry looks very similar to the overall breakdown, with ‘Improper Input Validation’ and ‘Buffer Overflow’ at the top of the list for most industries.

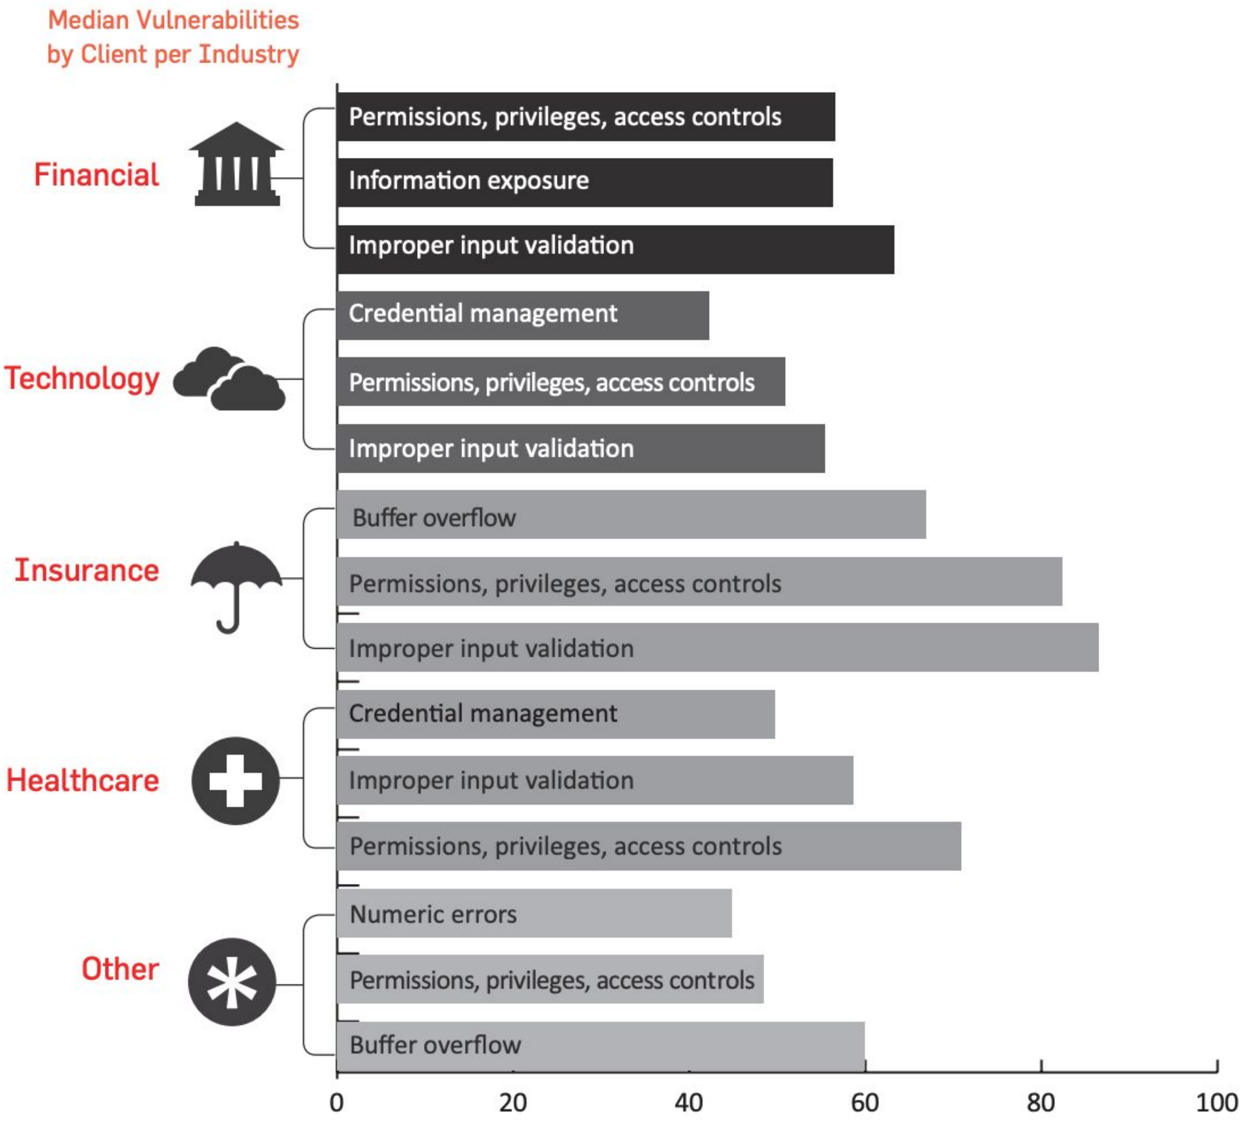

In the previous vendor analysis section, we identified a specific set of vulnerable vendors clients in each industry used. However, here we see that the top vulnerability weaknesses are more or less the same across industries. This seems to indicate that regardless of vendor/product, each industry deals with the same types of weaknesses; mainly related to lack of access controls and to improper input validation enabling XSS and SQL injection in web applications.

CVSS – Common Vulnerability Scoring System – Score Analysis

Our analysis of customer vulnerabilities points to a set of reasons why the CVSS score is a weak foundation for risk-driven automation. Most scanners prioritize vulnerabilities based on their CVSS score, which is based on six factors: Authentication, Access Vector, Access Complexity, Confidentiality Impact, Availability Impact, and Integrity Impact.

A majority of the vulnerabilities we find on our client's systems have the highest CVSS score of 10.0. In contrast, only a small subset of vulnerabilities are associated with known and publicly documented attacks. About 25% of vulnerabilities listed in the National Vulnerability Database (NVD) have public exploits on Exploit DB, and only around 3% are used in a specific form of malware (also known as exploit kits).

In effect, the CVSS score blurs the distinction between practical and theoretical risk. Relying exclusively on the CVSS score leads to a higher volume of 'critical' vulnerabilities to sort through – and less ability to effectively prioritize the highest risk vulnerabilities.

CVSS and NopSec Risk Score Comparison

The following is a chart representing the distribution of CVSS scores for vulnerabilities found on NopSec clients' systems. Many of the detected vulnerabilities in NopSec's customer environments are scored as a 10.0, the highest possible CVSS score. In light of the small number of vulnerabilities that are actually exploited, this indicates that the CVSS score is not sufficient for automated risk-driven vulnerability prioritization, leading to a large number of falsely prioritized vulnerabilities.

Distribution of Client Vulnerability CVSS Scores

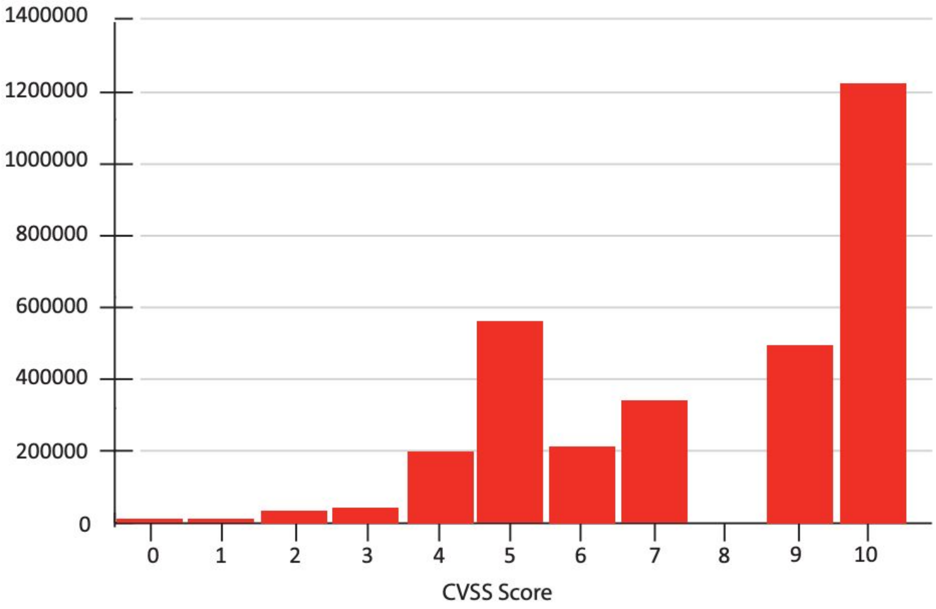

The graph below reinforces this analysis. It depicts the median CVSS base score for vulnerabilities found on client systems in each industry. Of particular interest is the Financial industry, where the median score is 9.3. This means that over half of the vulnerabilities found on financial clients' systems score a 9 or a 10 on the CVSS base score, making prioritization using only the CVSS score impossible.

Median CVSS Base Score

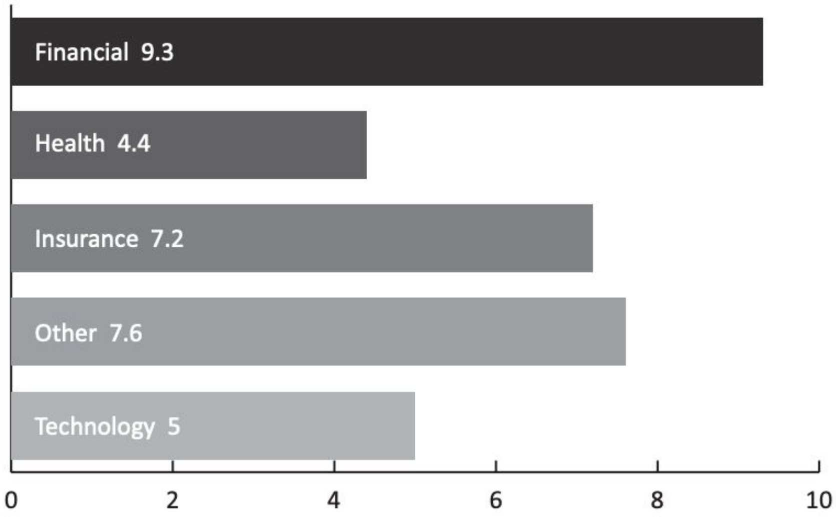

All vulnerabilities in NopSec's client database are also assigned a Technical Risk Score (calculated on a scale from 0-1). This proprietary risk calculation takes into account external factors including availability of a public exploit, malware correlation, and social media trends to evaluate the risk of each vulnerability within a broader context than the CVSS score alone.The histogram below depicts the distribution of the Technical Risk Score of these same client vulnerabilities. For greater granularity, we also plotted this data on a logarithmic scale.

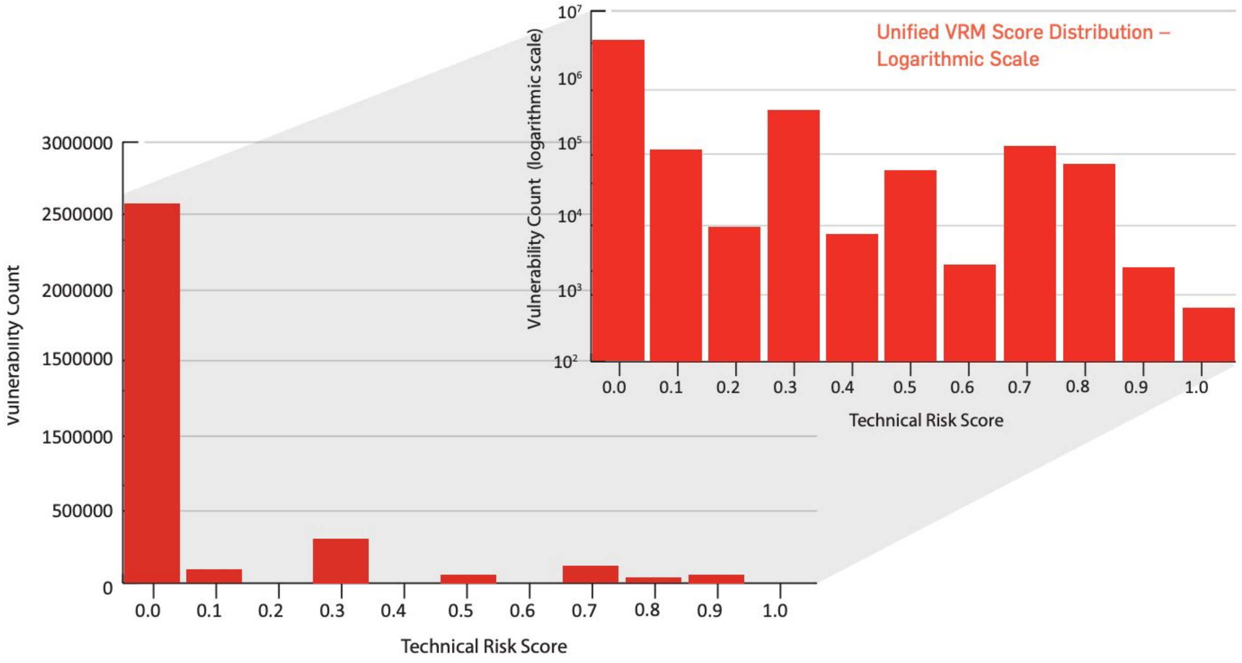

Unified VRM Score Distribution

After calculating the Technical Risk Score, most vulnerabilities are classified on the low end of the spectrum, with only a small subset assigned the highest scores. Taking into account a more multi-dimensional, contextual, and data-driven model has the outcome of generating a more manageable prioritization of vulnerabilities and risks.

Poor Predictive Value in Likelihood of Attack

Because the CVSS score is only based on the attributes of the vulnerability itself, and does not incorporate external information like public exploits, remediation programs will find it falls short as an effective foundation for prioritization. The section above focused on the “critical-heavy” aspect of CVSS, which could lead to many falsely prioritized vulnerabilities. This section will explore the shortcomings of CVSS in terms of predicting attacks, which leads to false negatives.

The CVSS score can be divided into two parts, or sub-scores: The Exploitability sub-score (measured using Authentication, Access Vector, and Access Complexity), and the Impact sub-score (measured using Confidentiality, Availability, and Integrity Impact). Examining the exploitability sub-score, our research found that the factors incorporated in the score do not have high value in predicting whether a vulnerability will actually be exploited. This could be because the majority of vulnerabilities have a local access vector and a lack of authentication resulting in a “flattening” of equivalence between vulnerabilities with real risk or danger of an exploit and “safer” vulnerabilities.

The distribution of the CVSS Exploitability sub-score confirms this, as the majority of vulnerabilities are rated with a 9 or higher.

Distribution of CVSS Exploit Score

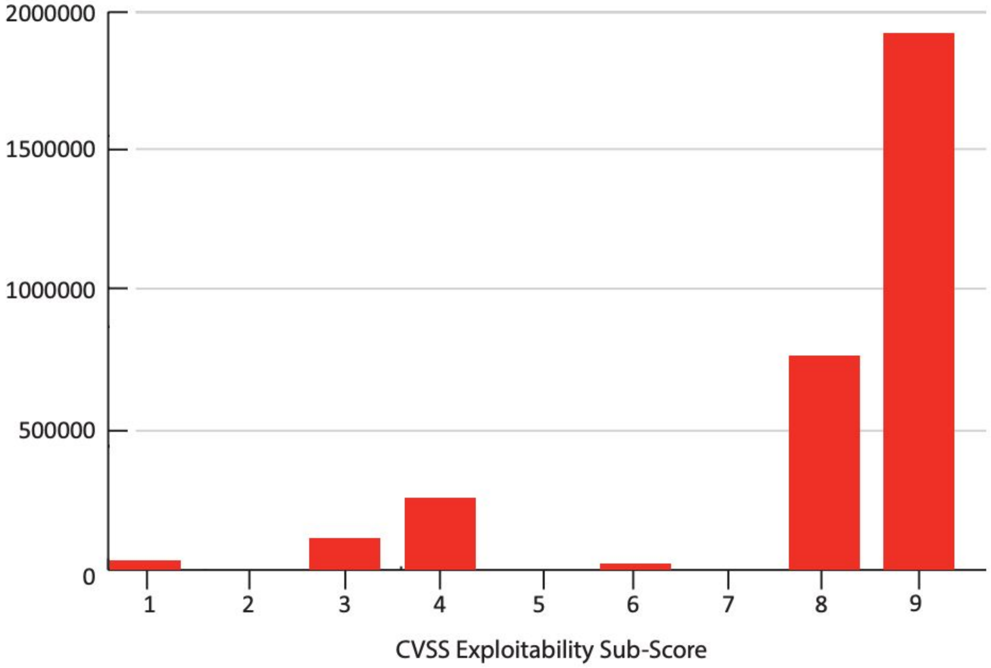

The median CVSS exploitability scores for clients in each industry further illustrates the problem. For clients in Financial and Technology industries, over half of the vulnerabilities present on their systems have a CVSS Exploitability sub-score of 10.0, the highest value. This stands in sharp contrast to the observation that only about a quarter of vulnerabilities have exploit code available, and that a significantly smaller percentage are actively exploited.

Continuing with our analysis of the CVSS exploitability sub-score, how do its components (access vector, access complexity, and authentication) differ for vulnerabilities with public exploits and for vulnerabilities used in malware? The results are below.

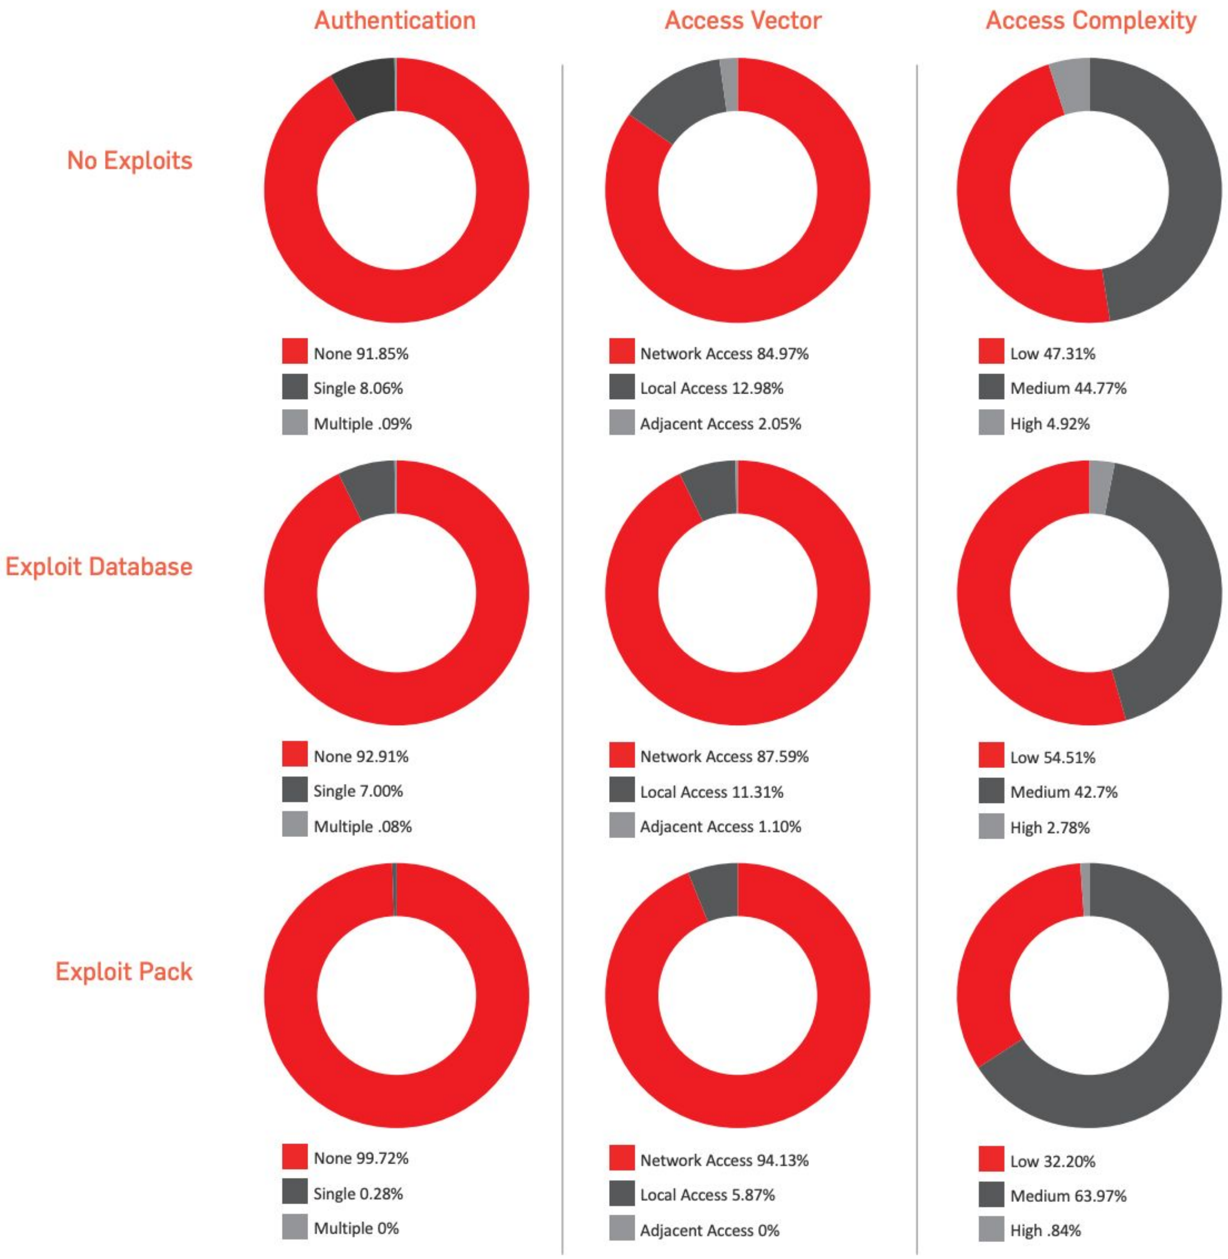

From these results, we see that the breakdown of each category does not differentiate well between vulnerabilities with public exploits and/or active malware. This further indicates that the CVSS Exploitability score is not a good predictor of the likelihood a vulnerability will be exploited.

Recognizing Value: CVSS Impact and Ease of Exploit

The other half of the CVSS score is the Impact sub-score, which attempts to measure the data loss that could occur if the vulnerability is exploited. Data loss is divided into three categories: availability loss, integrity loss, and confidentiality loss. Each is measured on a scale from None to Partial to Complete.

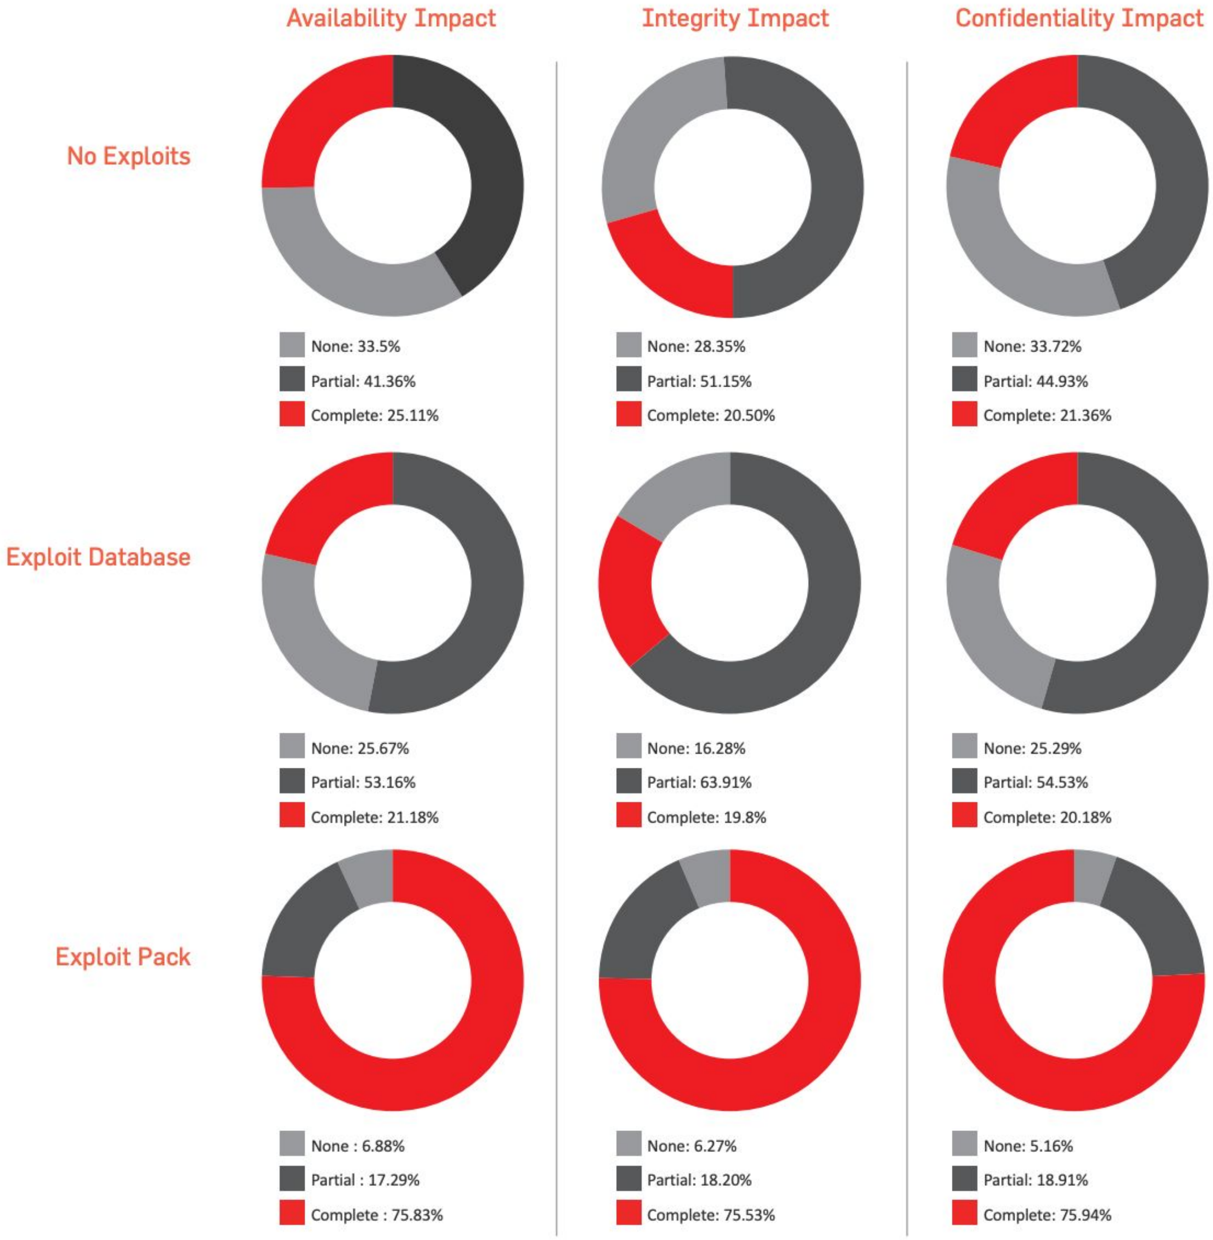

Do vulnerabilities that have public exploits and have active malware associated with them measure differently on these axes than vulnerabilities that have not?

Each of the CVSS Impact categories do seem to be correlated with whether or not a vulnerability is exploited. This would indicate that attackers care less about how easy a vulnerability is to exploit, and more about the actual impact and outcome of the exploited vulnerability. This, in turn, might affect the intrinsic value of the vulnerability and the monetary outcome of a vulnerability sale in the zero-day market. Even if a vulnerability is easy to exploit, it is not worth a hacker's time unless it will actually lead to a certain result of a compromised asset. This is consistent with our earlier observation that attackers are motivated not only by the relative ease with which they can exploit a vulnerability, but also the relative value of the asset where the vulnerability resides.

It's Who You Know: Improving Risk Analysis Through Social Media Correlation

Recent research indicates that social media, particularly Twitter, is becoming the go-to resource for security researchers and attackers looking to disseminate proof-of-concept exploits. In order to keep organizations up-to-date with the latest vulnerability trends, NopSec collects and incorporates Twitter data into its risk evaluation. As we highlighted in our 2015 State of Vulnerability Risk Management report, NopSec found a direct correlation between social media interactions and the risk a

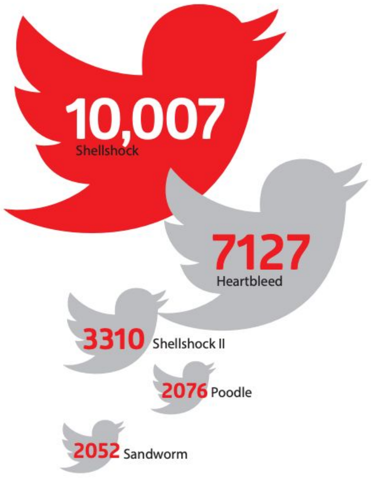

The separate instances of Shellshock are due to an improper initial patching, prompting a second Shellshock vulnerability

vulnerability poses to an organization. The top five most tweeted CVEs (as in our database) reflect Twitter interactions focused on well-publicized and dangerous vulnerabilities.

While some highly dangerous vulnerabilities have thousands of Twitter interactions, the majority of vulnerabilities are never tweeted about or are only tweeted about once. The graphic below presents the median number of tweets for all vulnerabilities, vulnerabilities with a public exploit, and vulnerabilities associated with active malware.

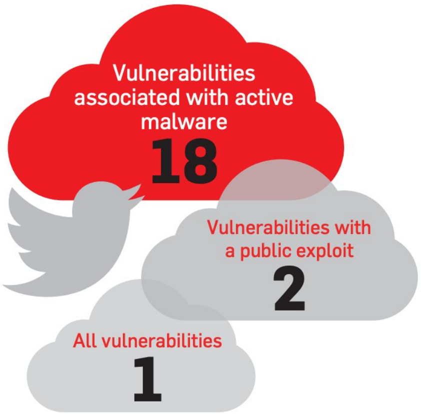

The large difference in median tweets between all vulnerabilities and vulnerabilities with active malware indicates that Twitter interactions are highly correlated with the danger a vulnerability presents to organizations. Additionally, the difference in median tweets between vulnerabilities with a public exploit and vulnerabilities with active malware shows that Twitter is an excellent differentiator between vulnerabilities that present only moderate risk (a public exploit is available, but the vulnerability may not have been exploited in the wild), and those that present significant risk (vulnerabilities being actively exploited in the wild).

Malware-Based Vulnerability Risk Evaluation

FireEye™

In evaluating whether or not a vulnerability represents a threat to the organization, NopSec considers, among other things, whether the vulnerabilities are used (“weaponized”) by active malware in the wild. In order to provide organizations with a more comprehensive view of risk, NopSec partnered with FireEye to analyze vulnerabilities currently exploited as part of malware or targeted hacking campaigns. As part of this analysis, NopSec correlated 54 vulnerability CVEs found by FireEye in recent exploits and attacks with anonymized client data to get a detailed view of the current threats faced by our clients.

Of the vulnerabilities, most of the them were found on client systems in the Financial Industry, followed by Technology. This is at least partially due to the large number of assets clients in the Financial industry tend to have. However, the high number of weaponized vulnerabilities present on clients in the Financial industry's systems does represent significant danger.

Industry Breakdown

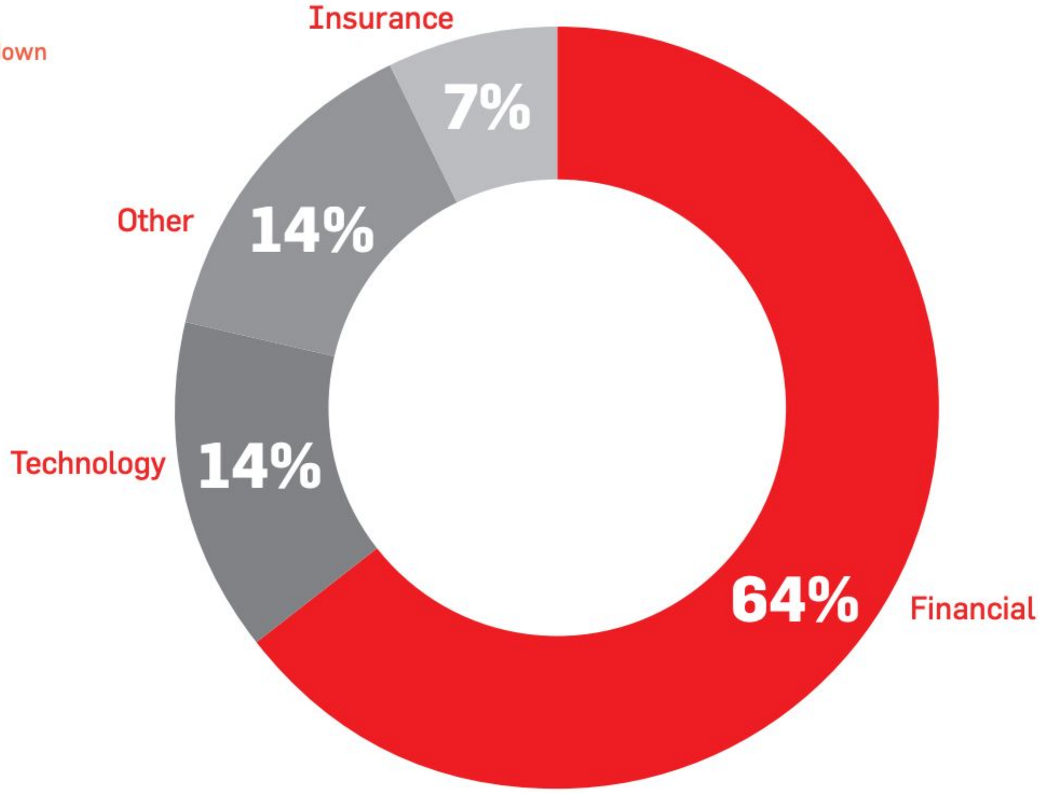

The exploits provided by FireEye are primarily vulnerabilities in Microsoft Windows Office, Microsoft Silverlight, Adobe Flash and Oracle/Sun Java. Given the wide use of all of these products, this indicates that these exploits present significant risk to clients across industries.

Vendor Breakdown

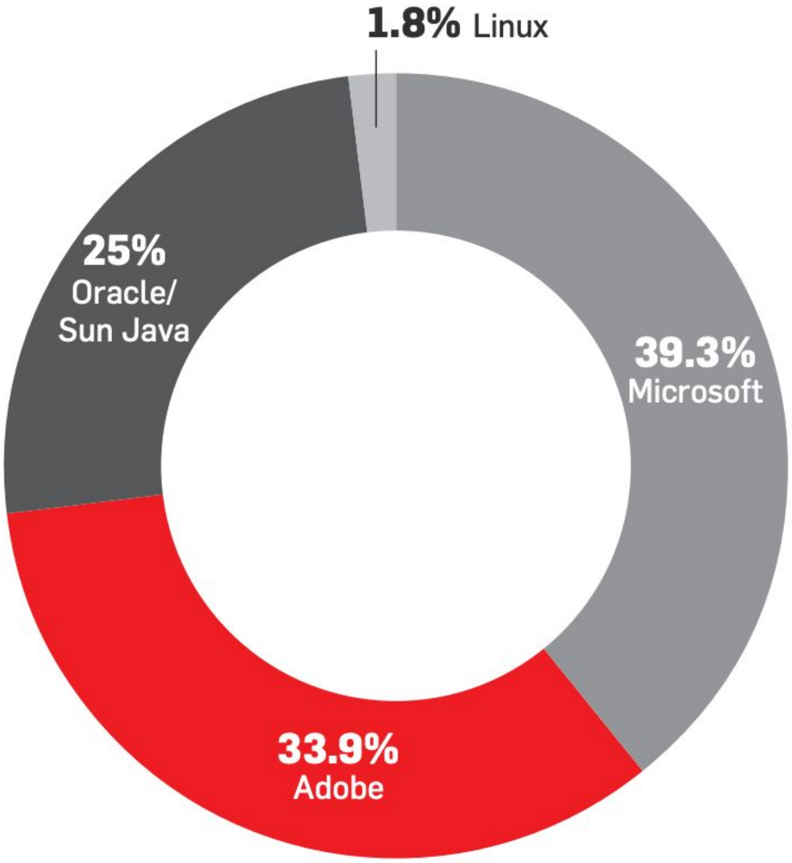

All assets affected by these vulnerabilities in the NopSec customer database were part of the customers' internal networks. This is particularly relevant because, as we will discuss later, customers often think that their internal assets are "safer" and therefore fail to quickly remediate vulnerabilities on these assets. Exploiting these vulnerabilities is trivial through spear phishing attacks that attempt to get targets to open attachments or visit controlled web pages containing embedded exploits.

Active Malware CVEs – Exploit Kit Breakdown

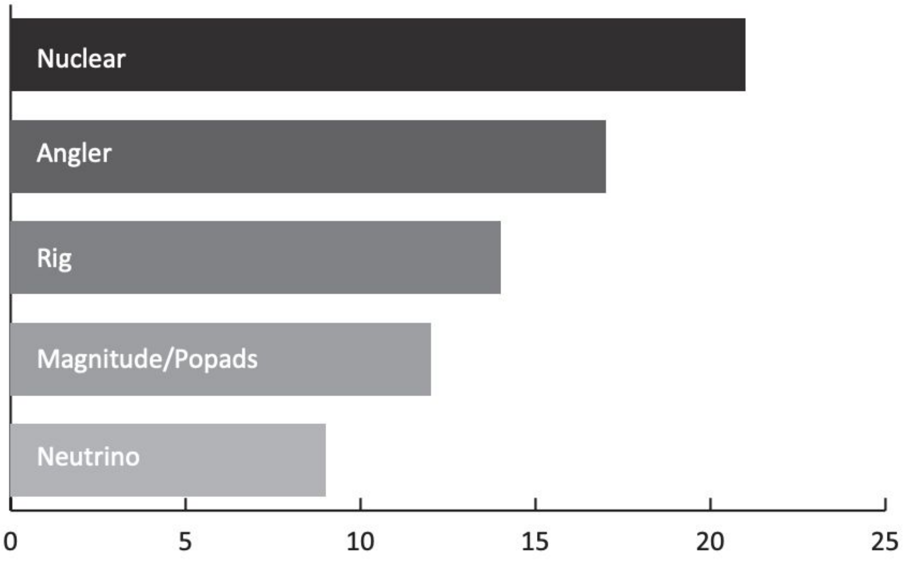

Many of the CVEs associated with active malware had a high amount of social media feed, with the median number of tweets for the CVEs provided by FireEye being 121.

FireEye™

The much higher Twitter mentions of exploits for CVEs we monitor reflects the fact that these vulnerabilities exist in major enterprise software and services. These software and services are lucrative targets for threat actors given their widespread use. Similarly, it follows that Nuclear and Angler, as some of the most sophisticated exploit kits on the market that are quickest to incorporate exploits for new vulnerabilities, would target the software and services most likely to exist on a potential victim's system. The focus on targeting systems in the financial industry may follow from a similar logic, as actors may target more *lucrative systems with access to financial data and credentials, though we do not have sufficient data to determine whether this information accurately reflects Angler's and Nuclear's usage, or is more a reflection of the demographics of NopSec's client data.*

Exploit Window Analysis

Given that time to remediation is a critical dimension in reducing risk and managing the attack surface, we also investigated the exploit window – the time between a vulnerability being confirmed and it being exploited. To approximate the exploit window, we calculated the time between the original publication date of a vulnerability in the National Vulnerability Database and the first exploit listed in the Exploit Database.

Many of the CVEs associated with active malware had a high amount of social media feed, with the median number of tweets for the CVEs provided by FireEye being 121.

There are two key considerations to this analysis – first, most exploits in the Exploit Database are published at the same time as or before the corresponding vulnerability is published in the National Vulnerability Database. This is because often a vulnerability will not be considered "confirmed" and therefore ready for publication until a working exploit is developed. Additionally, vendors will often wait to formally disclose vulnerabilities until they have developed a corresponding patch. While this data does show a need for improvements in the vulnerability disclosure process, it does not imply that most exploits are so-called "zero-days."

Second, it is important to distinguish between an exploit published in the Exploit Database and a vulnerability being exploited in the wild. Exploit disclosure is an important part of the vulnerability disclosure process, and public exploits are often published by security researchers as a Proof of Concept of exploitation of disclosed vulnerabilities to the vendor. This means that while public exploit availability is an important component of risk, it is not a guarantee that a vulnerability is being exploited in the wild. NopSec incorporates many indicators of exploitation in addition to presence in the Exploit Database, such as Metasploit module availability and active malware correlation, in order to evaluate vulnerability risk with as much detail as possible.

The graph below presents the number of vulnerabilities with exploits published over various windows of time.

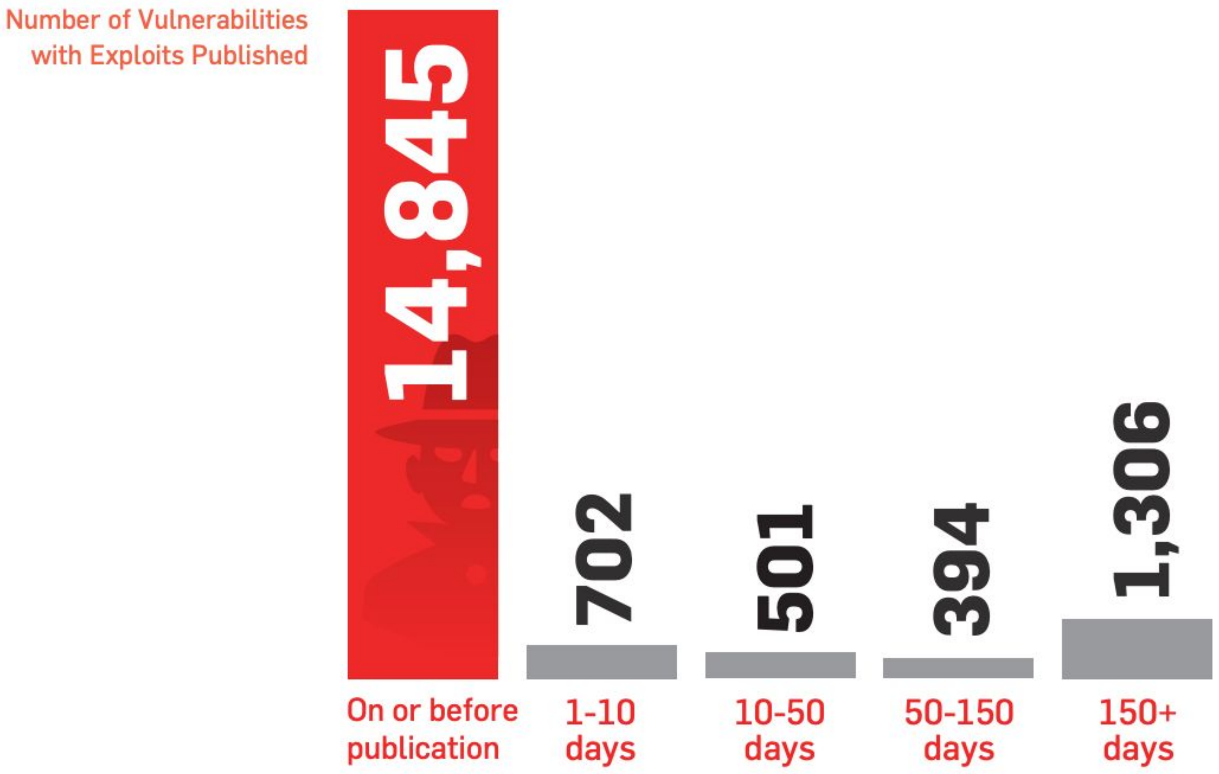

The median exploit window for all vulnerabilities is -5 (exploit published 5 days before CVE publication). However, narrowing in on vulnerabilities exploited only after publication, the median rises to 104 days. This means that even if a vulnerability does not have a corresponding public exploit code, a public exploit could be available in a few months. The wide variance in the exploit window illustrates the importance of an automated and consistent vulnerability remediation program that stays up-to-date with the latest exploit information. These conclusions are similar to a recent study by FireEye, which found that proof-of-concept exploits are on average published 17 days before manufacturers publish a patch. However, malware kits were slower to incorporate exploits, with one kit even incorporating an exploit 412 days after patch publication.

The wide variance in the exploit window illustrates the importance of an automated and consistent vulnerability remediation program that stays up-to-date with the latest exploit information.

Vulnerability Remediation

The central concept of NopSec's Unified VRM is vulnerability remediation. This is essential to reducing the overall security risk and exposure to an acceptable level over time. Building on our exploit window analysis, we can postulate that reducing the time to remediation — which spans vulnerability identification to successful mitigation – is crucial to improving the risk profile of an organization by reducing the potential attack surface.

Remediation Times by Atttack Vector (in Days)

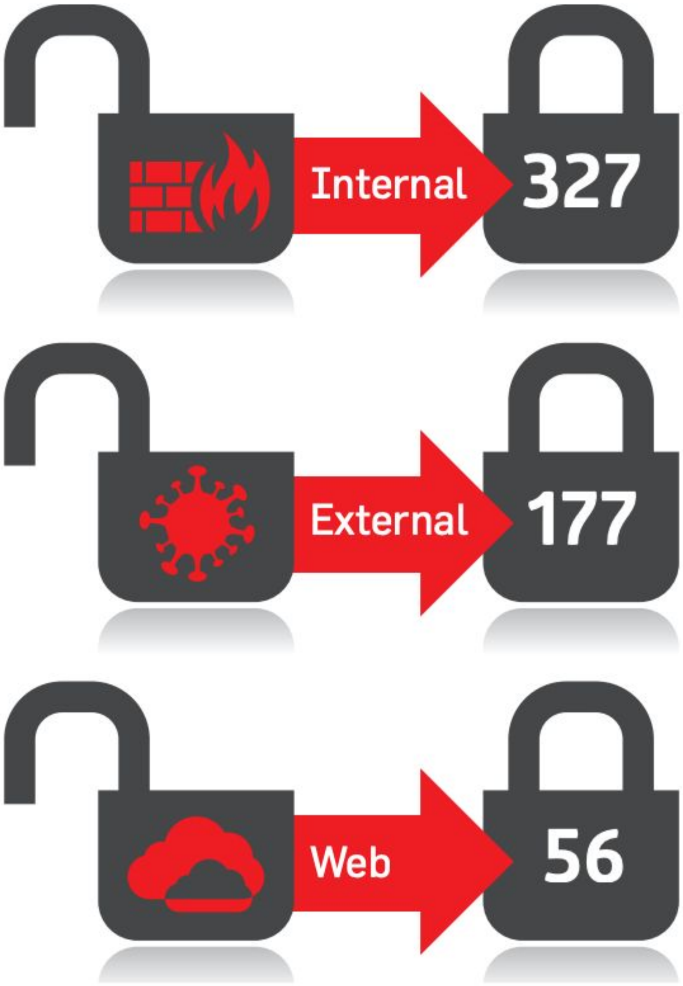

The analysis of remediation times by module/attack vector (external, internal, web) shows it takes longer for clients to remediate vulnerabilities present in internal networks since the team responsible for remediation likely assigns a lower risk to vulnerabilities discovered on internal assets. This is in sharp contrast to the increase of phishing attacks targeting internal networks that exploit known vulnerabilities through malware "drive-bys" or other methodologies. On the other hand, it takes less time to remediate web vulnerabilities, possibly because vulnerabilities on web applications are client-facing and therefore more exposed. Additionally, the developers that build web applications may be more sensitive to security vulnerability risks and secure coding practices, leading to better remediation times.

Conclusion

The analysis that emerges from our 2016 report is consistent with research done by other organizations such as FireEye : organizations struggle to accurately assess the risks posed by vulnerabilities, and to implement an efficient remediation program that tackles risk by reducing the attack surface. A new approach to a comprehensive and integrated enterprise vulnerability management program is needed to prioritize remedial actions on vulnerabilities that represent the most risk of exploitation and on assets that carry the most monetary value in terms of protected data.

Vulnerability remediation is central to the success of an organization's vulnerability risk management program. However, organizations delay their vulnerability remediation regardless of the vulnerability criticality.

Our analysis reinforces the conclusion that the CVSS Base is an imperfect measure of vulnerability risk and a poor predictor of tangible threat through malware that exploits the vulnerability. Narrowing down and improving filtering of vulnerability scanning output based on criticality, risk, and asset value are the foundation to improving the remediation process. Relying on the CVSS score is not only misleading – but can also prove to be counterproductive.

Instead, the analysis supports our view that the components of the NopSec risk score are good standalone and combined predictors of vulnerability risk that include:

- Presence of a vulnerability public exploit as a measure of the likelihood of attack
- Availability of malware and documented instances of exploits leveraging the malware
- Correlation of known vulnerabilities with active malware and social media interactions

Security tools that prioritize threat prediction and likelihood of attack offer many benefits to organizations struggling with what to do next after detecting the multitude of security vulnerabilities across the IT environment. If organizations know what to focus on, the window of exposure and risk of a data breach is greatly reduced.

Find out how NopSec's Unified VRM can help you think like a hacker and stay ahead of the trends. Visit www.nopsec.com or email hello@nopsec.com for additional information or to request a demo.

About NopSec Labs

The research contained in this report was conducted by NopSec Labs. All NopSec data presented in this report is compared with the population data from National Vulnerability Database, where relevant. NopSec has a leading research and data science practice focused on analyzing malware, exploit, vulnerability, and other cyber threat risk patterns. Our team of data scientists applies that knowledge to help organizations forecast the probability of a data breach and improve prioritization, remediation, and reporting of critical vulnerabilities. Customers of NopSec's Unified VRM platform are provided with a variety of reports specific to their organization and similar to the data contained in this report.

NopSec helps security professionals simplify their work, effectively manage and prioritize vulnerabilities, and make better informed decisions.

NopSec's Unified VRM is an innovative threat and vulnerability management solution that addresses the need for better prioritization and remediation of security vulnerabilities in a single platform.

NopSec Inc. • www.nopsec.com • info@nopsec.com