KENNA Security

Cyentia INSTITUTE

PRIORITIZATION TO PREDICTION

Volume 3: Winning the Remediation Race

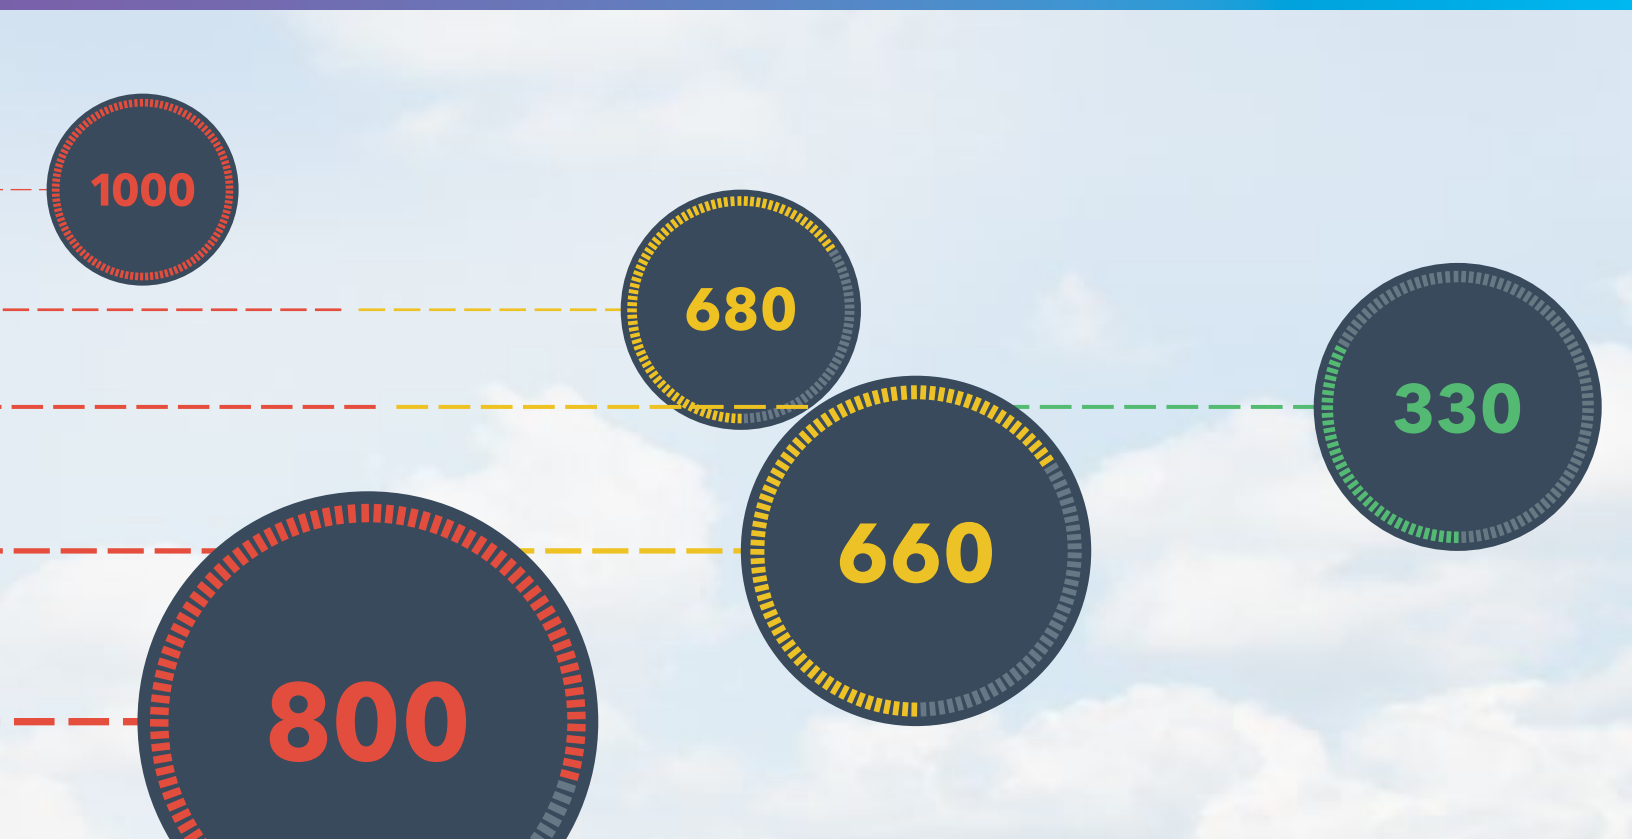

PRIORITIZATION TO PREDICTION

VOLUME 3: WINNING THE REMEDIATION RACE

KENNA Security

This research was commissioned by Kenna Security. Kenna collected and provided the datasets for this reserarch to the Cyentia Institute for independent analysis and drafting of this report.

Kenna Security is a leader in predictive cyber risk. The Kenna Security Platform enables organizations to work cross-functionally to determine and remediate cyber risks. Kenna leverages Cyber Risk Context Technology™ to track and predict real-world exploitations, focusing security teams on what matters most. Headquartered in San Francisco, Kenna counts among its customers many Fortune 100 companies, and serves nearly every major vertical. For more information, visit www.kennasecurity.com.

119 Cyentia INSTITUTE

Analysis for this report was provided by the Cyentia Institute. Cyentia seeks to advance cybersecurity knowledge and practice through data-driven research. We curate knowledge for the community, partner with vendors to create analytical reports like this one, and help enterprises gain insight from their data.

Find out more: www.cyentia.com.

Introduction

> *Though I concede thy quicker parts, Things are not always done by starts, You may deride my awkward pace, But slow and steady wins the race.*
>
> — THE HARE AND TORTOISE BY ROBERT LLOYD

What is a pace that wins the race in vulnerability management? Do proverbial hares and tortoises exist—firms that possess "quicker parts" that aid them in that race? Is it possible to be both fast and steady? How is progress measured along the way? What does winning look like? We seek answers to these questions and more in this edition of Prioritization to Prediction.

The Prioritization to Prediction series is an ongoing research initiative between Kenna Security and the Cyentia Institute. The first volume proposed a model for predicting which of the numerous hardware and software vulnerabilities published each month were most likely to be exploited, and thus deserving of priority remediation. The second volume sought to apply and test that theoretical model using empirical data collected on billions of observed vulnerabilities. We ended the last report by analyzing vulnerability remediation timeframes across a sample of 12 firms.

This third volume picks up where we left off and expands the analysis to roughly 300 organizations of different types and sizes. We leverage a technique called survival analysis to draw out important lessons about remediation velocity and capacity, concepts we explore and define during the course of this report. Overall, our goal is to understand what it means to survive—nay thrive—in the race of vulnerability remediation.

Some key terms used in this report

We've made an effort to use conventional terminology in this report, but some things we discuss don't have a clear or universally accepted term. In such cases, we chose a term that seemed appropriate and endeavored to use it consistently. This page should help keep us honest and avoid misinterpretation.

Exploited vulnerabilities:

This refers to either an *exploit* (proof-of-concept or working code for exploiting a vulnerability) and *exploitation* (attacks targeting a vulnerability in the wild). Since exploited vulnerabilities represent increased risk, we also refer to these as "high-risk" vulnerabilities.

Open/Closed vulnerabilities:

Vulnerabilities observed in an environment are in an open or closed state. Closed means the vulnerability has been patched, fixed, or otherwise remediated. Open means it has not been closed and thus exposes the affected asset(s) to any exploits that exist now or in the future.

Key Findings

- The median time-to-remediation across all firms is 100 days. 25% of vulnerabilities remain open over a year.
- Organizations close high-risk vulnerabilities almost twice as fast as they do others. Intelligence matters!
- The median remediation time for Microsoft vulnerabilities is 15X shorter than for Oracle, HP, and IBM. The lesson to vendors: make it easy for customers to fix your mistakes.
- Smaller firms tend to fix issues faster than their medium and large counterparts. Complexity appears to trump capacity in the long run.
- Remediation velocity differs across industries. Healthcare institutions, for example, take 5 times longer than leading industries to close vulnerabilities.
- The Banking, Oil/Gas/Energy, and Insurance sectors all outperform the mean for remediation coverage.
- Any given organization, regardless of size, can address about one out of every 10 vulnerabilities in its environment.
- Top-performing organizations remediate over twice the number of vulnerabilities at a rate three times faster than the norm.
- Most firms cannot (or barely) keep up with the rate of exploited vulnerabilities added to their environment each month. But one-third manage to gain ground in that race. Want to know what sets them apart? Keep reading!

Key Findings from Volumes 1 & 2

23% of published vulnerabilities have associated exploit code.

2% of published vulnerabilities have observed exploits in the wild.

50% of exploits publish within two weeks surrounding new vulnerabilities.

About one-third of published CVEs are actually observed in live enterprise environments.

Only 5% of all CVEs are both observed within organizations AND known to be exploited.

Some CVEs (about 3%) were observed across a million or more assets each.

40% of vulnerabilities observed in enterprise networks are still open (unremediated) today.

Dataset & Demographics

A sanitized sample of nearly 300 organizations was selected for inclusion in this study. Together the organizations observed over 2 billion vulnerabilities across their respective assets. This represents a sample of Kenna Security customers that met certain conditions allowing us to perform meaningful survival analysis on vulnerability remediation data. These conditions include minimum thresholds for the number of vulnerabilities, assets in scope, level of activity, length of time using the platform, etc. Kenna sanitized all data that could be used to identify customers before the Cyentia Institute's analysis.

FIGURE 1

Industries represented by percent of organizations and number of vulnerabilities

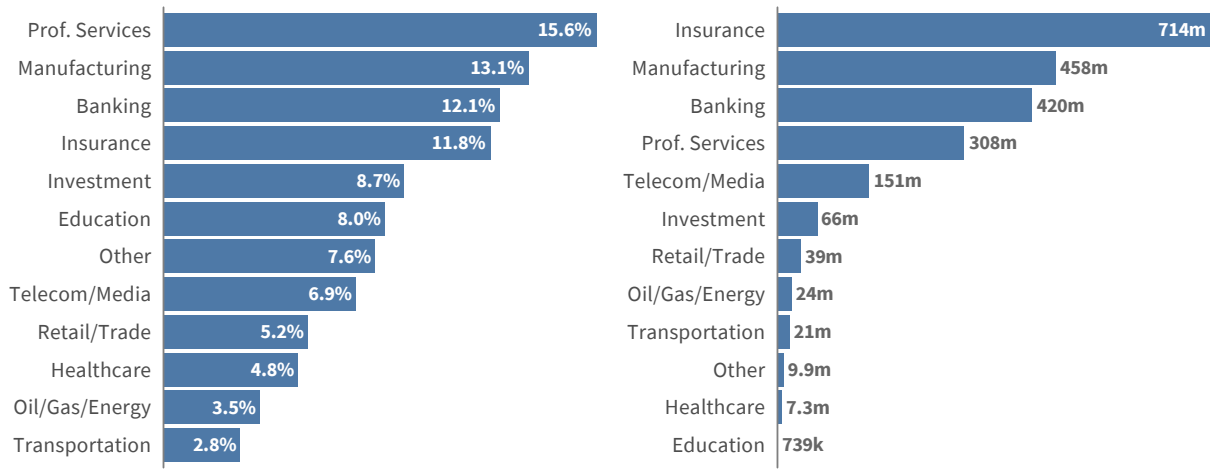

Source: Kenna / Cyentia

Figures 1 and 2 provide demographic information on the organizations included in this analysis. The first offers two views of the industries represented, one based on the percent of firms and the other based on the total number of vulnerabilities. We do this because the contribution to the corpus of vulnerabilities is not equal among all organizations or industries.

Figure 2 reveals a fairly even distribution across small, medium, and large firms. It also gives a range of employee counts for each of those categories to make it easy to determine where your organization fits. We will show findings broken out by the industries and sizes shown here at various points in this report.

FIGURE 2

Organization sizes represented based on number of employees.

Source: Kenna / Cyentia

Lessons in Remediation Velocity

As mentioned in the introduction, we ended the last report by analyzing vulnerability remediation timeframes across a sample of 12 firms. Here we extend that to nearly 300 organizations.

Survival analysis is a set of methods to understand the time duration to an event. In our sphere, the event of interest is the remediation of a vulnerability, and it's a very useful way to view timelines in vulnerability management. Here's a brief walkthrough of the basic principle and then we'll hop into the results. Let's assume an organization observes 100 live/open vulnerabilities within its assets today (Day Zero) and manages to fix 10 of them, leaving 90 to live another day. The survivability rate on Day Zero would be 90% with a 10% remediation rate. As time passes and vulnerabilities continue to be fixed, that proportion will continue to change.

Tracking this change over time produces a curve shown in Figure 3, where we see that, on average, it takes firms about a month to remediate 25% of vulnerabilities in their environment. Another two months get them over the halfway mark, pegging the median lifespan of a vulnerability at 100 days. Beyond that, there's clearly a long tail challenge for remediation programs that results in 25% of vulnerabilities remaining open after one year. The stairsteps or wiggles you see in the line between those points reflect the fits and starts of pushing patches.

FIGURE 3

Overall vulnerability survival analysis across firms.

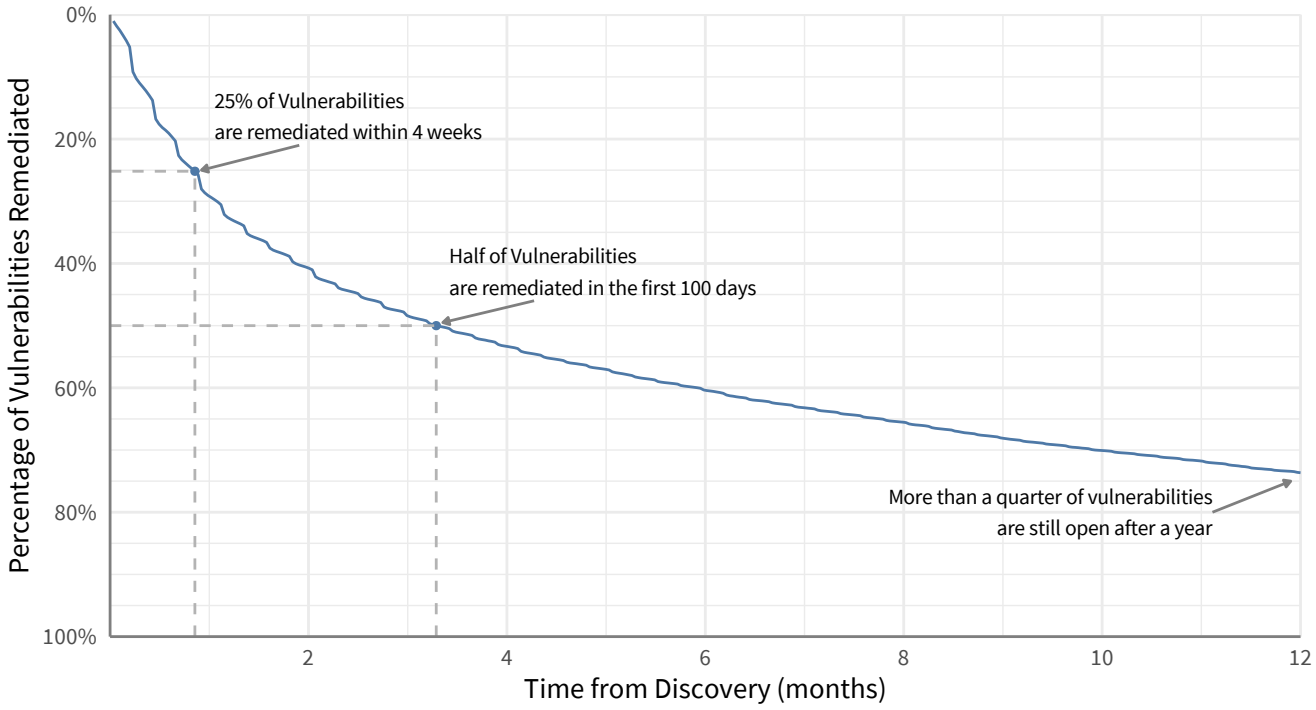

Source: Kenna / Cyentia

Figure 4 offers a simplified view of the data in Figure 3 that focuses on the first, second, and third quartiles. It typically takes 26 days for firms to remediate 25% of vulnerabilities, 100 for 50%, and 392 days to reach 75%. Many of the following charts adopt this view of survival analysis results to conserve space and convey results quickly.

FIGURE 4

Alt. view of overall vulnerability survival analysis across firms.

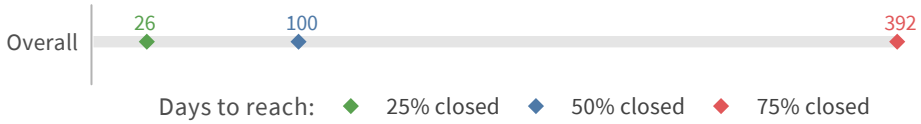

Source: Kenna / Cyentia

It should be intuitive that wide variation exists around the aggregate remediation timeline from Figure 3. We include Figure 5—which is admittedly an indecipherable mess—just to illustrate how wide that variation actually is among the organizations in our sample. It's plain that they cover the entire spectrum. We chose to color the lines by industry to further make the point. There is no obvious pattern to indicate, for instance, that all firms in a sector cluster together. That's not to say no differences exist among industries at the aggregate level; they do and we'll cover that a bit later.

FIGURE 5

Individual vulnerability survival analysis for each firm. You are not supposed to be able to read this.

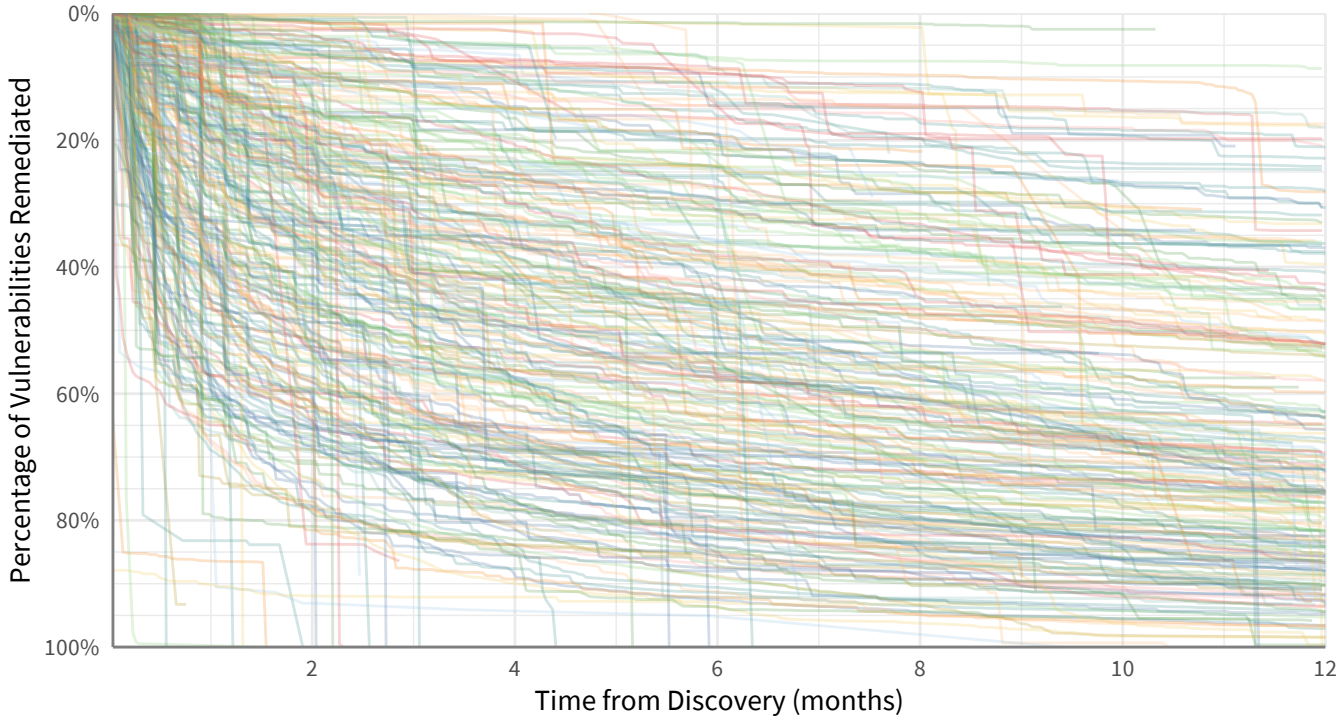

Source: Kenna / Cyentia

A review of Figure 5 begs the question of why such extreme variations exist among firms in terms of remediation timelines, especially since they all have the common denominator of being Kenna Security customers. But beyond that one commonality, consider the vast range of factors inside and outside these organizations that might impact these results. From that perspective, perhaps these differences aren't so surprising at all. Much of this third volume focuses on understanding what some of those factors are and measuring their influence on vulnerability management programs.

Introducing Remediation Velocity

Before proceeding further, we need to define a new concept—remediation velocity. In survival analysis, the area under the curve (AUC) is an important measure. A lower AUC points to a shorter survival rate; higher means the opposite. For our application, that translates to a shorter lifespan for vulnerabilities. So the AUC provides a way to measure the velocity (speed and direction) of remediation efforts.

FIGURE 6

Example AUC measures for four firms.

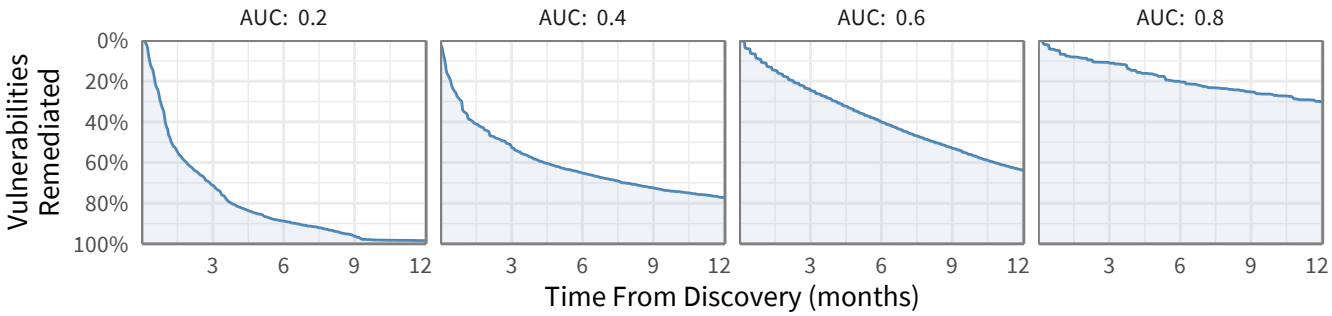

Source: Kenna / Cyentia

To illustrate this concept, we include Figure 6. It shows survival analysis curves for four organizations chosen to represent a range of survival curves. It's easy to see the relationship between the shaded portion under each curve and the resulting AUC. It's also easy to see how the firm on the left drives down vulnerabilities with more expediency than the others shown. We refer to this vulnerability survival timeline as *remediation velocity*.

PATCHING, FAST AND SLOW

When we presented the survival analysis for 12 example firms in the last report, some made the astute observation that it's not necessarily a bad thing that organizations do not close all vulnerabilities quickly. That is absolutely correct. In fact, it would be downright wasteful to drive to zero as fast as possible on every vulnerability.

As evidence of that, firms with high remediation velocity tend to have lower efficiency[1] ratings. This indicates a tradeoff between remediation velocity and efficiency that is similar to the tradeoff discussed in prior reports between coverage[2] and efficiency. This makes total sense and is a perfectly acceptable tradeoff to make for many organizations. Especially when the cost of not remediating (or doing it too slowly) is potentially much higher than over-remediating.

The more important question is how quickly a firm fixes the vulnerabilities that really matter. We will explore more about what that means later. For now, just know that's why we've chosen to examine coverage through the lens of velocity and capacity in this report.

---

1 *Efficiency measures the precision of remediation efforts. Of all vulnerabilities identified for remediation, what percentage should have been remediated?*

2 *Coverage measures the completeness of remediation efforts. Of all vulnerabilities that should be remediated, what percentage was correctly remediated?*

Comparing Vendors

A major external factor influencing the variation in remediation velocity ties back to the vendors that created those vulnerabilities in the first place. While it is true that remediation is an internal process, there's a lot happening outside the control of vulnerability management programs that impact timelines in substantial ways.

On that note, Figure 7 paints an almost unbelievable picture of the remediation timelines associated with major product vendors. It takes 15 times longer for firms to address half their vulnerabilities affecting Oracle, HP, and IBM products than to reach that same milestone with Microsoft products! What's more, moving from 50% to 75% remediation takes multiple years for several vendors.

FIGURE 7:
Remediation velocity for major product vendors.

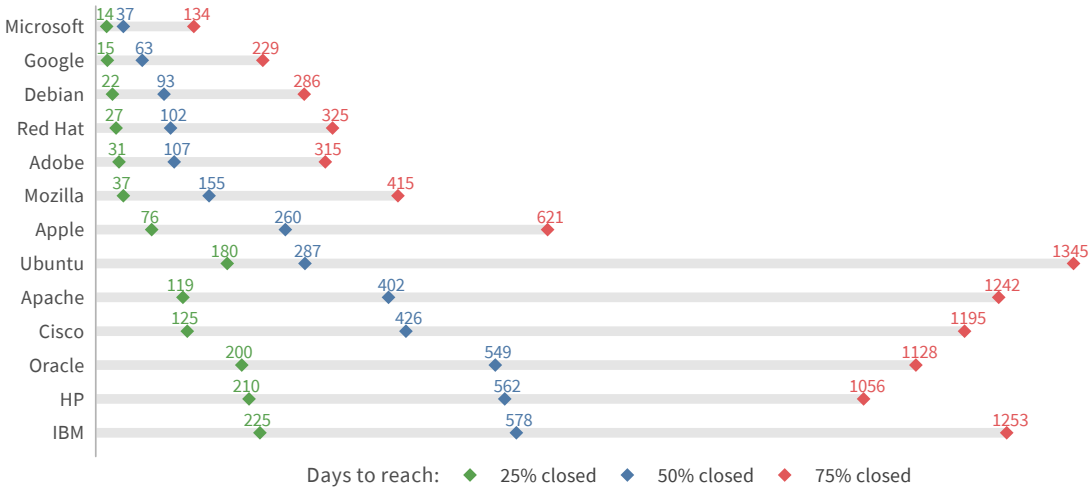

Source: Kenna / Cyentia

There are many possible reasons behind these dramatic differences. Java (Oracle) is notoriously hard to fix without breaking something. Apple might have fewer vulnerabilities than Microsoft, but their enterprise management support lags. Google updates frequently, but many forget to restart their browsers. Overall, though, we view Figure 7 as strong evidence in favor of scheduled releases and the automated distribution of patches. The lesson to vendors: make it easy for customers to fix your mistakes.

FIGURE 8
Remediation velocity for major product vendors.

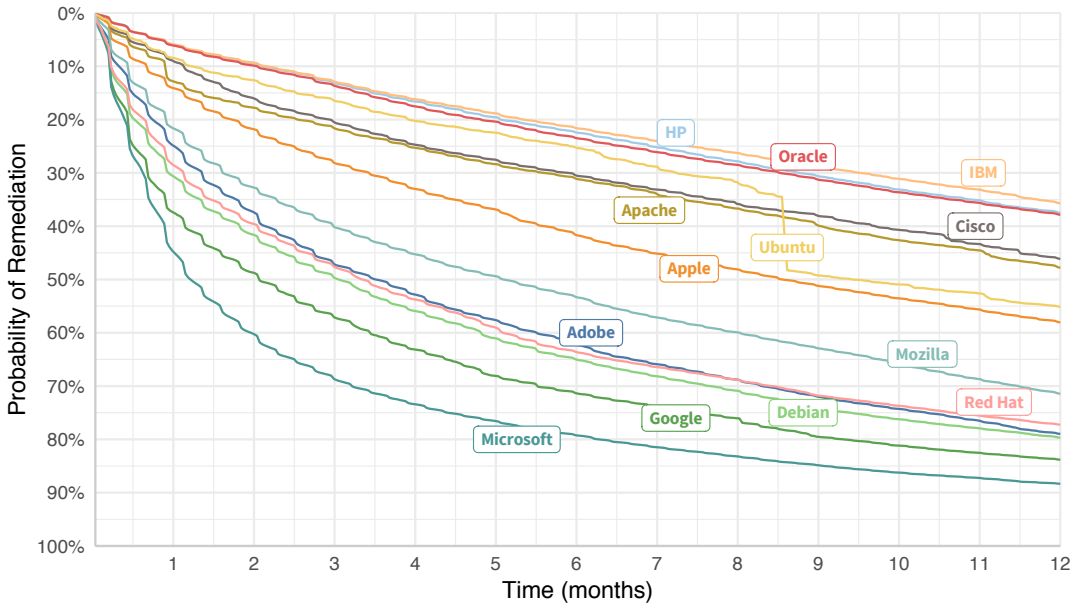

Apologies for the gratuitous data visualization, but we simply couldn't abide the thought of not including Figure 8. It gives an alternate view of the same results from the preceding figure. In addition to being pretty, it has a functional value of calibrating comparisons between the two types of charts.

Comparing Severity

The severity of a vulnerability represents a sort of transition between external and internal factors that may affect remediation velocity. The external component stems from the fact that severity metrics like the Common Vulnerability Scoring System (CVSS) are typically assigned by other entities. But there's an internal component as well because each organization decides how to act upon this information.

Intuitively, we'd expect to see organizations remediating high-severity vulnerabilities faster than those with low-severity ratings. Oh, that vulnerability management were that easy! But alas, the data delivers a firm, yet respectful "No" to that line of reasoning. Figure 9 asserts pretty much the opposite, in fact. Vulnerabilities with a severity score of zero to three (the lowest in CVSS) get fixed the fastest, while the 10s exhibit among the slowest remediation timeline.

FIGURE 9:

Remediation velocity by CVSS score.

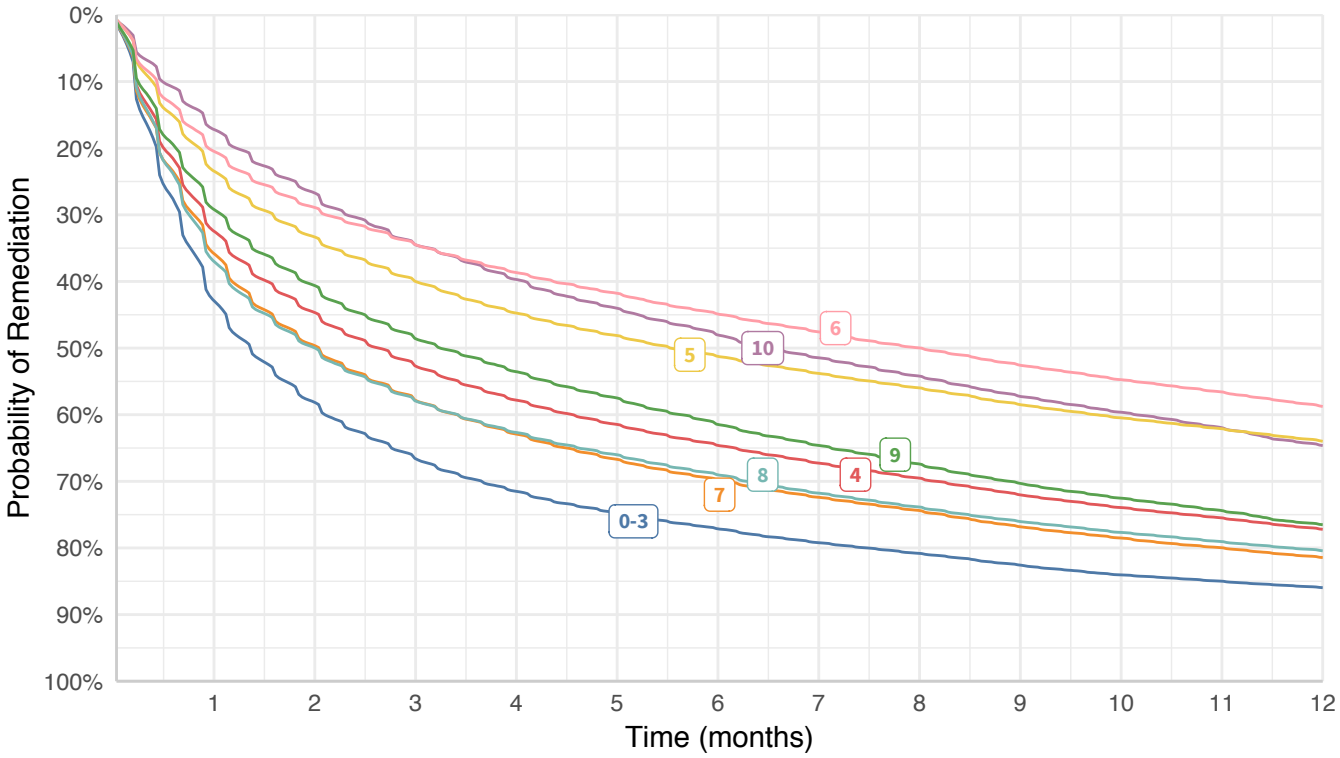

What gives? Well, either CVSS scores don't actually reflect severity, firms choose to ignore them, or firms cannot do anything with this information. The first two possibilities are difficult to test, but we do have a way to look into the third.

Comparing Exploited Vulnerabilities

Thankfully, we have the means of testing whether organizations have the desire and ability to alter remediation activities based on external intelligence about which vulnerabilities warrant attention. Organizations in our sample get that intelligence in the form of recommendations from the Kenna Security Platform. That intelligence includes data on which vulnerabilities are associated with previously published exploit code or exploit codes in the wild. We showed in prior reports that prioritizing these "easily exploitable" vulnerabilities outperforms other rule-based strategies, including CVSS.

We can perform the survival analysis based on just these easily exploitable vulnerabilities and compare the results to those with no known exploits. Here, we'd expect to see evidence that firms go after the exploited subset with more expediency. Figure 10 confirms this to be the case. Organizations close 25% of exploited vulnerabilities a week faster and reaching the halfway point takes roughly half as long. Over the long term, the delta between the groups at the 75% remediation milestone is measured in hundreds of days.

FIGURE 10:

Remediation velocity for exploited vs. non-exploited vulnerabilities.

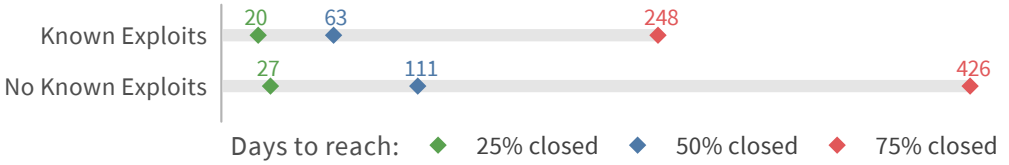

Source: Kenna / Cyentia

This is certainly good news for Kenna customers, but there's an important lesson for everyone else. The fact that organizations don't seem to be prioritizing remediation based on severity as defined by CVSS is not because they lack the desire or ability to do so. That suggests remediation velocity can be informed and improved by external intelligence about which vulnerabilities warrant attention. Thus, we conclude that the CVSS conundrum from Figure 9 has more to do with utility than ability.

Comparing Industry Sectors

Of course, the success or failure of a vulnerability management program doesn't hang on which products are in use or how external entities assess severity or exploitability. A host of internal factors play a role as well. We don't have information on the culture, governance, and security operations of the firms in our sample (though we hope to gain that insight for future research). But we do have demographic information like industry and number of employees, which we analyze in the following sections.

Comparing industries is always interesting because they often share similar business models, regulatory requirements, resource levels, etc. We've already learned from Figure 5 that firms within the same industry vary widely when it comes to remediation velocity, but the question in view here is whether any higher-order trends exist. Figures 11 (all vulnerabilities) and 12 (exploited vulnerabilities) offer perspectives on that question, but we recommend focusing on the big takeaways rather than the small differences.

We went back and forth about whether it was more appropriate to show survival analysis results for all vulnerabilities or just those classified as easily exploitable. On the one hand, the exploited subset is a better measure of remediation speed where it really counts. But on the other hand, using all vulnerabilities is a better measure of overall remediation capacity (how many can be closed in a given timeframe). So we've taken the tact of showing the view we find most relevant and interesting—and sometimes that means both.

FIGURE 11:

Remediation velocity by industry (all vulnerabilities).

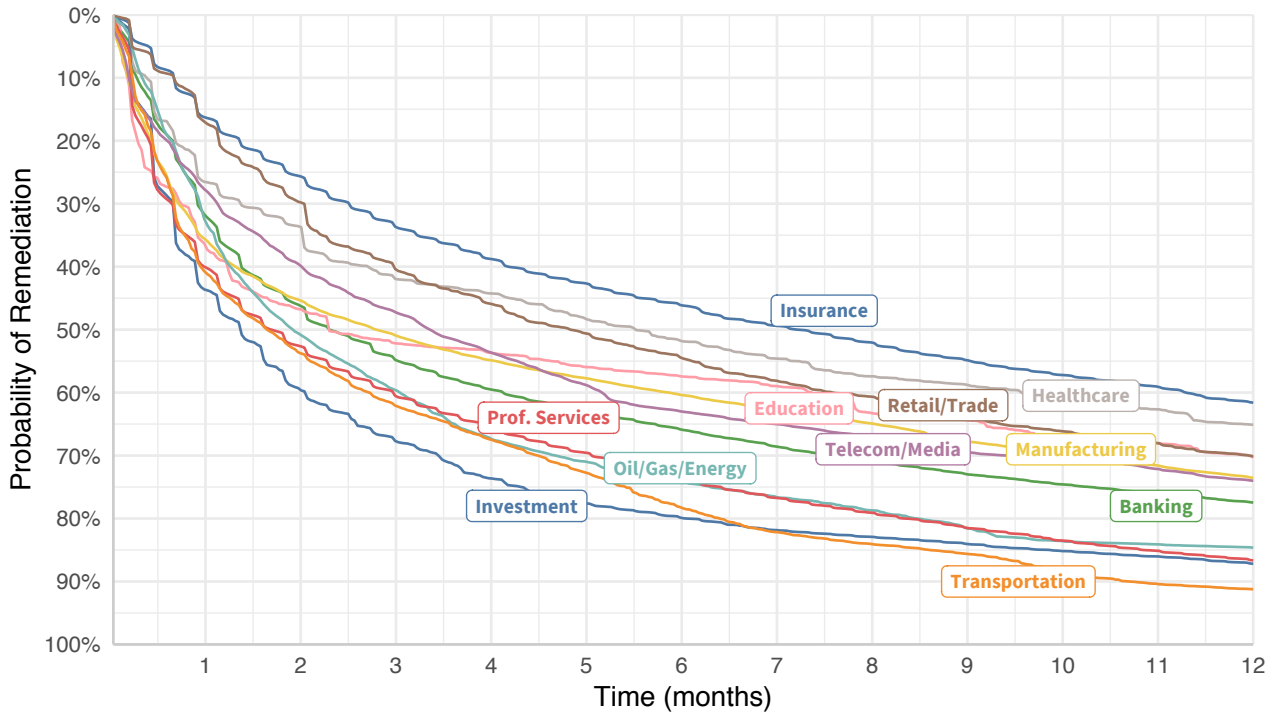

Financial Services (as a whole) and Healthcare are often contrasted as being polar opposites when it comes to cybersecurity posture. And true to the stereotype, we see Healthcare posting slow remediation velocity in both figures. Healthcare institutions take about 5 times longer than leading industries to close half their vulnerabilities.

Proving that stereotypes are never perfect, however, we see Banking in the middle and Insurance dead last (both subsectors of Financial Services). It's hard to fathom educational institutions remediating vulnerabilities faster than banks, but that's exactly what the data shows here. At least to the 25% milestone. Those lines diverge quickly after the midway point in Banking's favor, suggesting that being quick out of the gate isn't necessarily the best way to win the race.

FIGURE 12:

Remediation velocity by industry (exploited vulnerabilities).

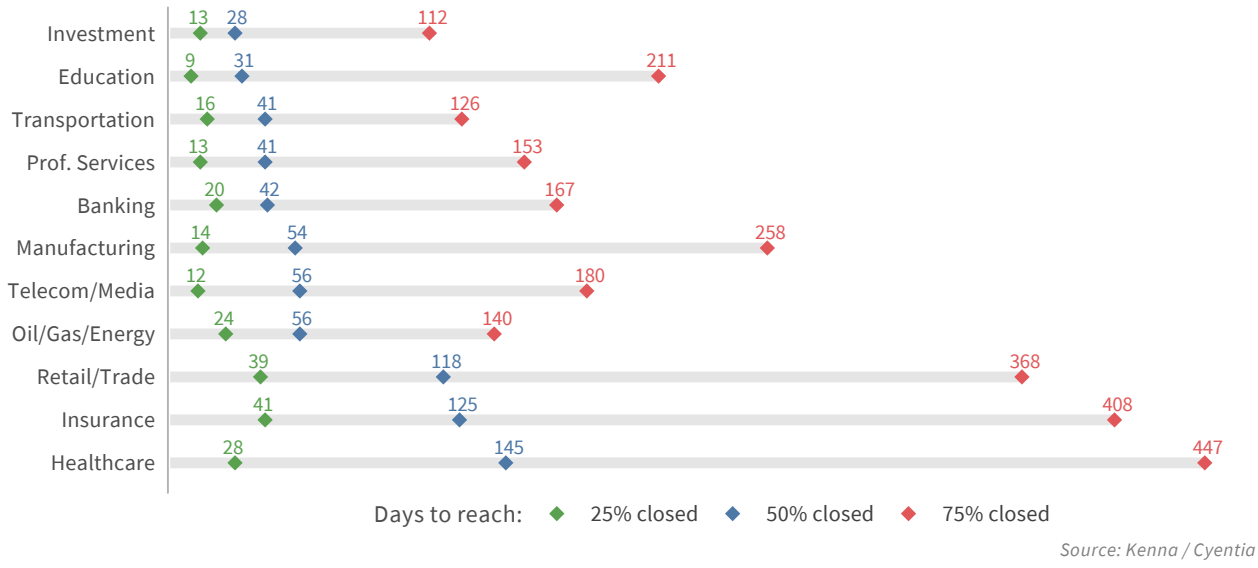

Source: Kenna / Cyentia

Speaking of starting gates and finish lines, it's interesting to note from Figure 12 that some industries have comparatively short interquartile ranges[4] for remediating exploited vulnerabilities. The range for the top three industries is nearly 3 times that of the bottom three. In a world where getting the job done in a reasonable timeframe is more important than getting started quickly, this may be the goal to shoot for.

To follow up on that point, we wanted to add a coverage-based view of remediation by industry. Recall that coverage measures the completeness of remediation. A remediation strategy that maximizes coverage over efficiency could be considered risk-averse, and that makes sense given what we see in Figure 13. The Banking, Oil/Gas/Energy, and Insurance sectors all outperform the mean for remediation coverage. Healthcare flatlines at the bottom again. Grab the defibrillator—stat! Curiously, the Investment sector posts negative gains from the top of Figure 12 to near the bottom of Figure 13 (they're quick, but less thorough).

FIGURE 13:

Relative remediation coverage by industry.

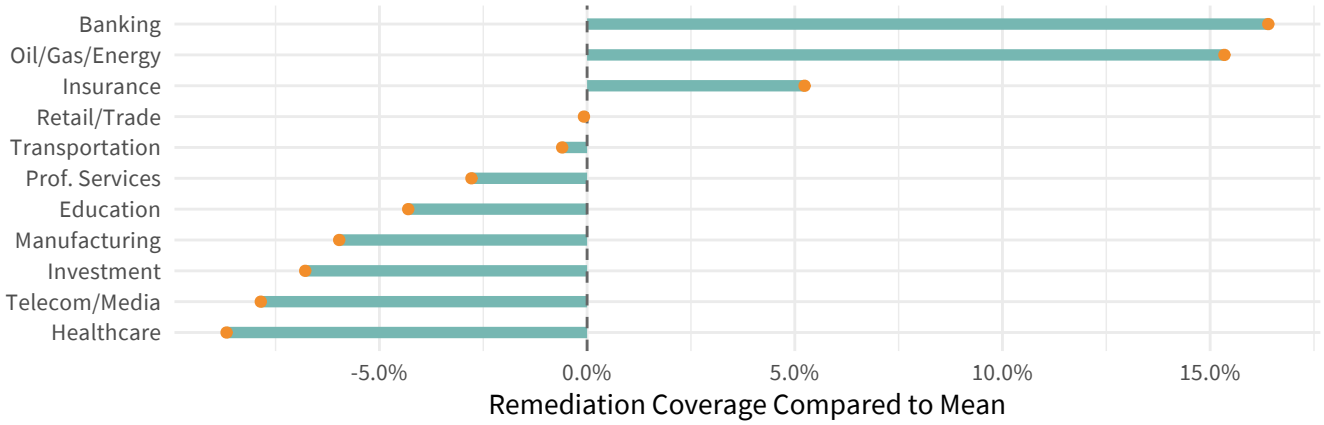

Source: Kenna / Cyentia

If you're thinking firms should aim for high marks in both coverage and remediation velocity, hold that thought. You'll need it again soon.

Comparing Organization Sizes

Another internal factor that may affect remediation velocity is organization size, and we suspect it's somewhat of a double-edged sword. Smaller firms have fewer resources, but they also tend to have smaller-scale problems. Larger firms get the luxury of more resources but not without more problems to boot. In that sense, remediation velocity may be a self-normalizing metric across firm sizes. But let's quit speculating and see what the data says.

FIGURE 14:

Remediation velocity by organization size.

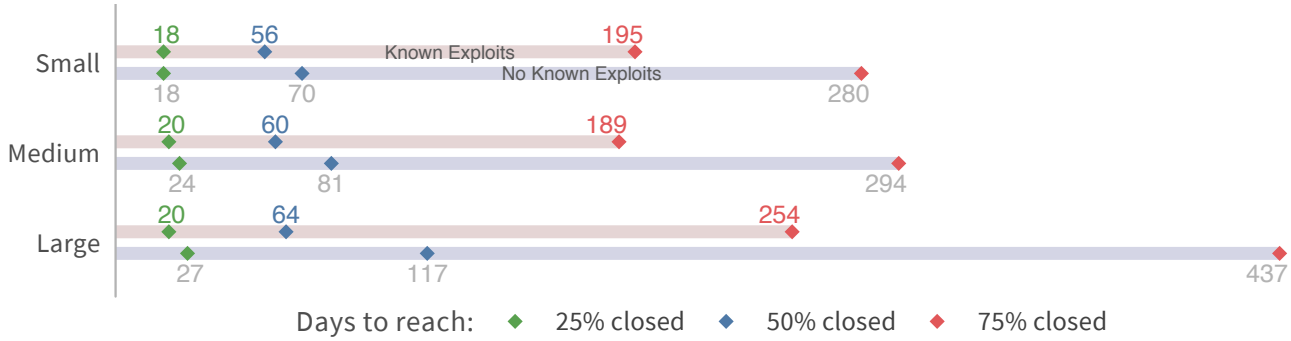

Source: Kenna / Cyentia

We suspect Figure 14 may challenge some preconceived notions. While many assume fewer resources would translate to reduced capacity to remediate vulnerabilities, smaller firms generally reach each time-to-fix milestone faster than their medium and large counterparts. And that trend is even more apparent when looking at exploited vulnerabilities.

Another aspect of Figure 14 deserving mention concerns the time interval between 50% and 75% remediation. First, we see noticeably shorter survival times for the exploited vulnerabilities. We want those stamped out quickly, so that's good. The second observation is more subtle but has important implications. Notice how the 50% to 75% interval for non-exploited vulnerabilities in large organizations is extremely long. It's almost like they've accepted they don't have the capacity to fix everything and have shifted resources accordingly to get more bang for the buck on those riskier vulnerabilities. And that's not a bad strategy.

In truth, Figure 14 probably says less about remediation capacity and more about the compounding difficulty of managing larger IT environments. The lesson we take from that is complexity trumps capacity in the long run. We're going to pick up this complexity-capacity thread again, but we need to hit one more topic before doing that.

Who Wins the Race?

Results thus far have hinted that raw speed might not be as important to winning the remediation race as the stamina to finish. But having both of those attributes would definitely be an advantage. Let's see if we can identify any biathletes among our sample firms.

Figure 15 removes any distinction of size or industry and simply looks at two key measures of remediation performance—velocity and coverage. Each dot represents an organization and its position on the Coverage axis corresponds to the proportion of exploited vulnerabilities remediated. The vertical axis shows where they rank in terms of remediation velocity. Firms highlighted in the top right can be considered top performers because they manage to address most of their high-risk issues and do so very quickly. #Winning

FIGURE 15:

Remediation coverage vs. velocity.

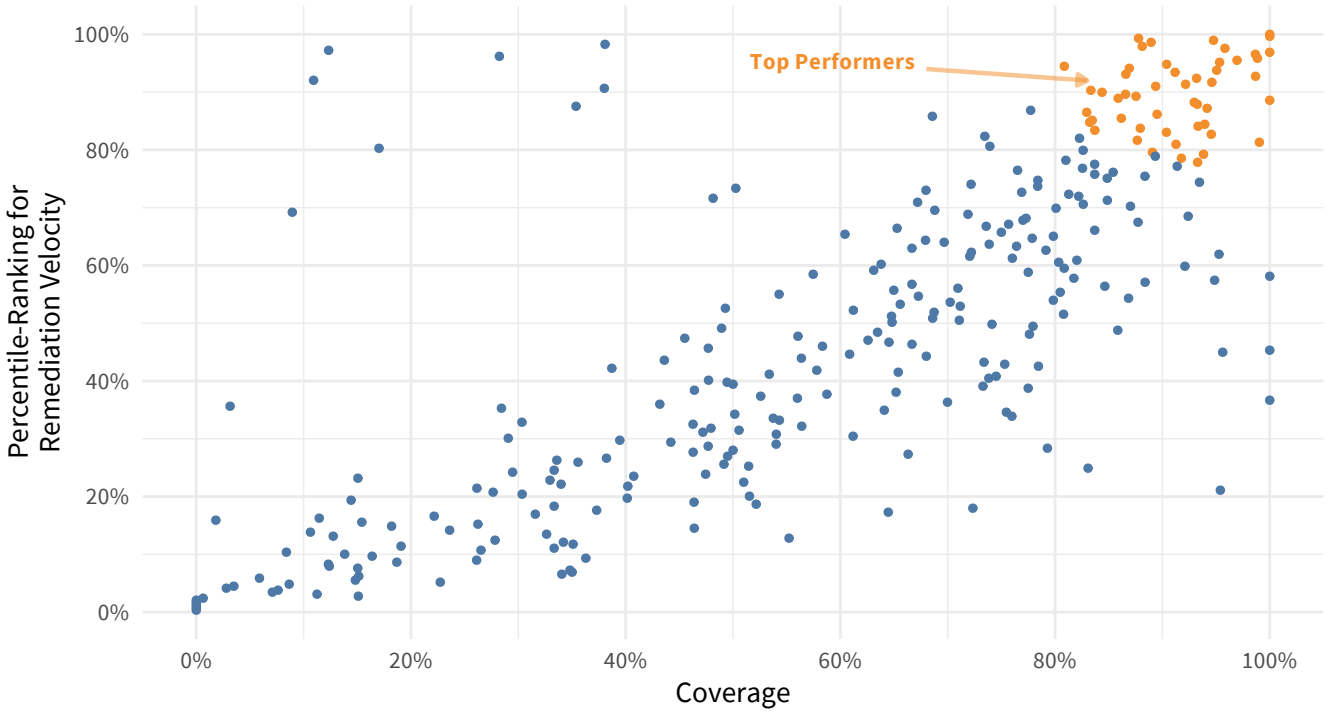

Source: Kenna / Cyentia

If you're a bit skeptical about these top performers, Figure 16 should convert you into a believer. It compares the remediation velocity of the top-performing firms from Figure 15 to the rest of the field. The difference is striking. The top performers start strong, stay strong, and remediate the majority of their vulnerabilities 3X faster than the others.

FIGURE 16:

Remediation velocity for top performers from Figure 15.

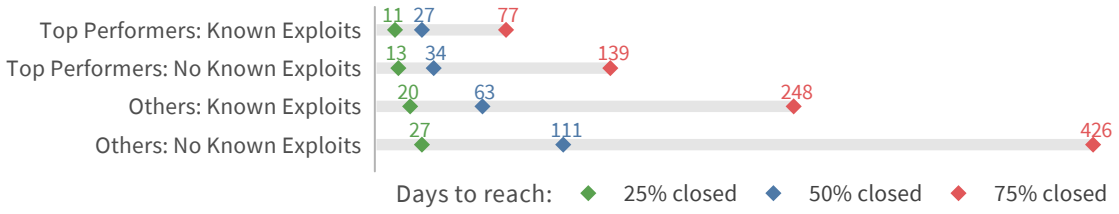

Source: Kenna / Cyentia

Perhaps we can glean some insight into these top performers by learning a little more about what types of organizations are among them. Figure 17 supplies an industry-based view on that, and Figure 18 reveals the size of those firms.

FIGURE 17:

Top-performing firms based on industry.

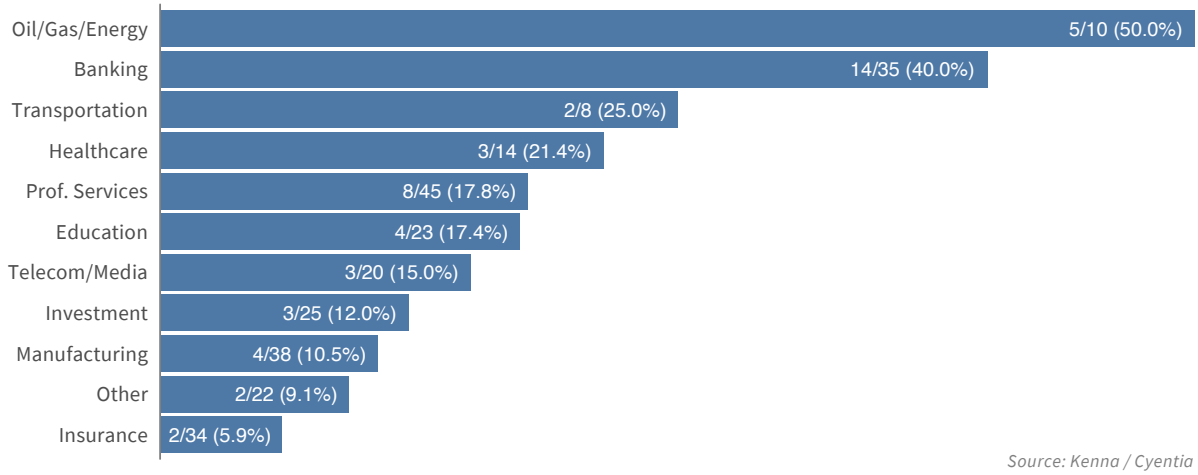

Source: Kenna / Cyentia

Overall, Figure 17 reflects findings similar to those we saw in Figure 13. Banking represents the largest contingent by far from a raw numbers perspective—14 of the 50 top-performing organizations. But those 14 banks represent 40% of all banks in our dataset, bumping that industry into a tie for the silver medal for relative proportion. They should still be proud, though, considering they maintained that high standard across a much larger sample size. Half of the firms in the Oil/Gas/Energy sector made the cut, giving it the pole position on the top performers list. Small sample sizes dominate the rest of the standings and make it difficult to give anyone clear bragging rights based on industry alone.

Referring way back to Figure 2, the breakdown of organization sizes across our sample is relatively even. But Figure 18 clearly shows a larger proportion of small and medium firms among the top performers. This finding aligns with the complexity-capacity tension we discussed earlier.

FIGURE 18:

Top-performing firms based on organization size.

Source: Kenna / Cyentia

The obvious question now is what, beyond simple demographics, separates these top performers from the pack? We already know they possess the winning combination of speed and stamina, but how did they acquire those traits? We attribute their edge to something we've termed remediation capacity. The next section shares what we've learned about that important characteristic.

Lessons in Remediation Capacity

We've learned a lot about remediation velocity, and in so doing, discovered some other concepts we need to better understand—asset complexity and remediation capacity. These concepts are admittedly difficult to define and measure, but we're going to give it a shot.

We would like asset complexity to measure the scope, diversity, intricacy, etc. of the IT environment. But we don't have access to all of that information, and so we will use the total monthly average number of vulnerabilities as a proxy for complexity. Our logic (which is backed by the data) is that more complex environments will have more assets with more software/services running on them, which inevitably means more vulnerabilities.

Remediation capacity measures the proportion of open vulnerabilities a firm can close within a given timeframe. To derive the remediation capacity of the firms in our sample, we first calculated their average number of closed and open vulnerabilities per month. We then used those to construct a regression model. Figure 19 records the results, which are, frankly, rather amazing.

FIGURE 19:
Ratio of open to closed vulnerabilities per month.

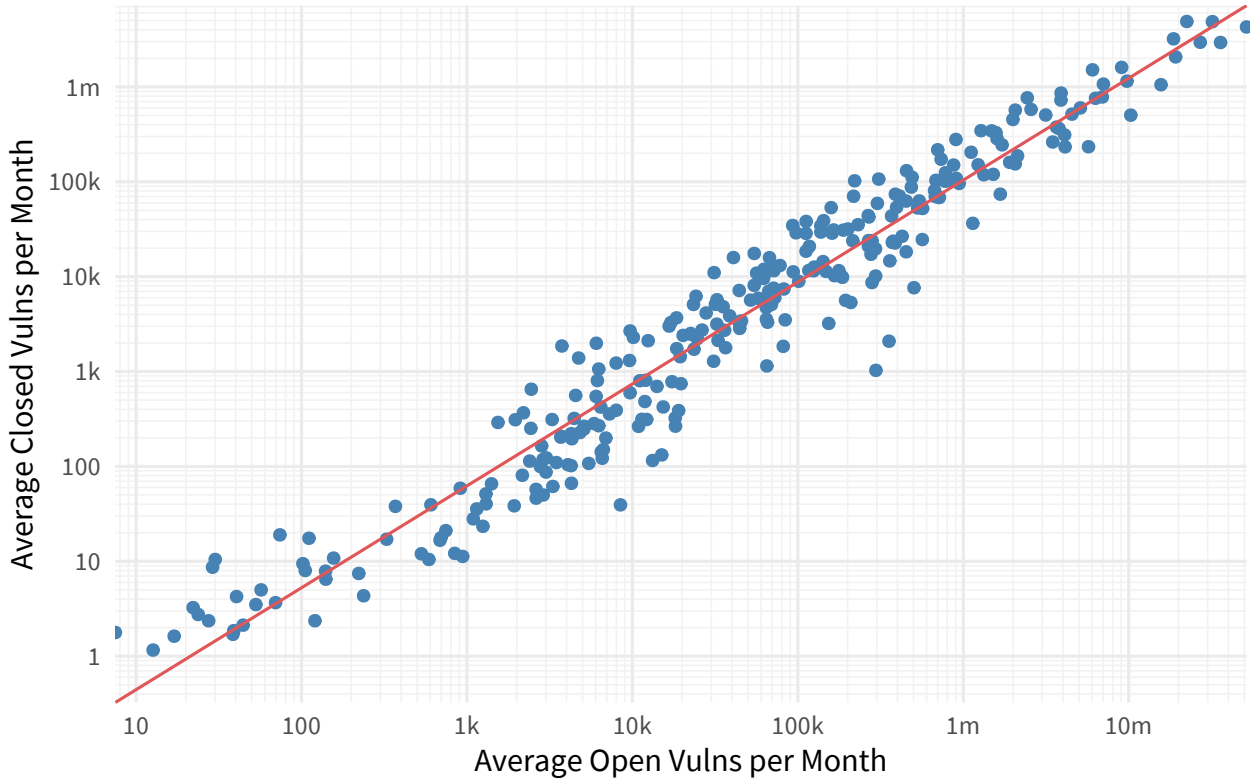

Source: Kenna / Cyentia

The R^2 statistic for this log-log regression model is 0.93, meaning that it's very strong and captures most of the variability around vulnerability closure rates. You can see this visually in Figure 19 because all the points—which represent the remediation capacity for each organization—fit tightly along the predicted trendline.

Strong models are great (especially when they're so simple), but there's something else Figure 19 teaches us that's greater still. Notice first that each axis is presented on a log scale, increasing by multiples of 10. Now, follow the regression line from the bottom left to upper right. See how every tenfold increase in open vulnerabilities is met with a roughly tenfold increase in closed vulnerabilities?

That, in a nutshell, is why it feels like your vulnerability management program can never pull ahead in the race of remediation. A typical organization, regardless of asset complexity, will have the capacity to remediate about one out of every 10 vulnerabilities in their environment within a given month. That seems to hold true for firms large, small, and anywhere in between. And it reaffirms that "more money, more problems" principle we discussed earlier.

So is there no hope? Are vulnerability remediation programs destined to slowly drown in a quagmire of their own making? No! Figure 20 offers a ray of hope and way out of the stalemate.

FIGURE 20:

Ratio of open to closed vulnerabilities per month for top performers vs. other firms.

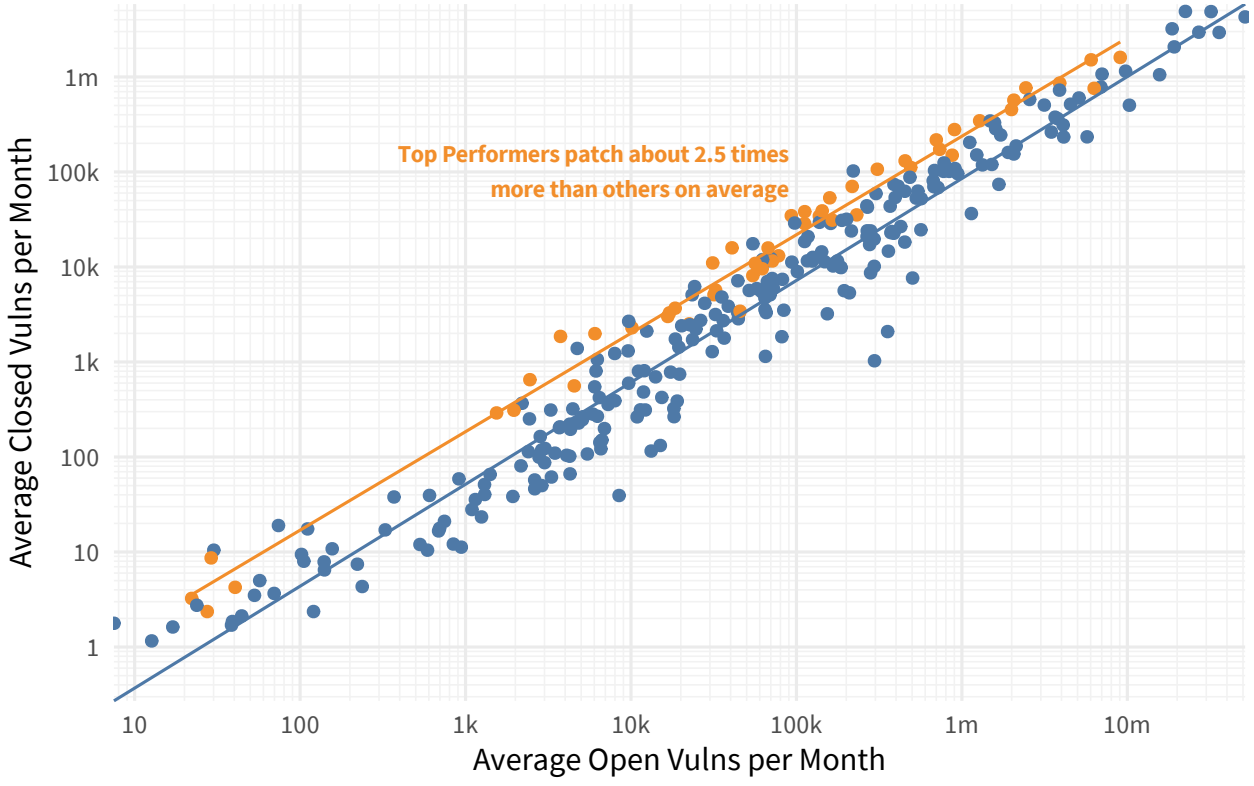

Source: Kenna / Cyentia

Figure 20 builds upon the previous figure by adding a second regression model for the top performers (shown in orange). That both the line and points (firms) fall above the original is wonderful news for weary programs. It means those top performers achieve a higher remediation capacity than other organizations. And that means the glass ceiling of remediating one out of every 10 vulnerabilities per month we established before can be broken. Shattered, in fact. Top performers exhibit two-and-a-half times the remediation capacity of other firms!

FIGURE 21:

Comparison of remediation capacity among top performing and other firms.

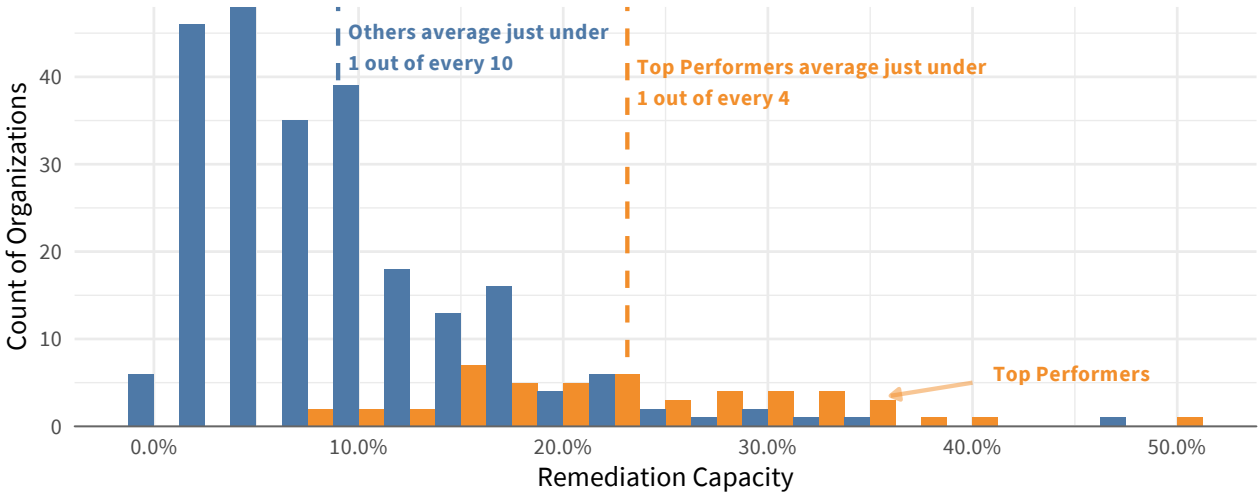

Source: Kenna / Cyentia

Figure 21 displays the comparison of remediation capacity between top performers and other firms in a different, perhaps more familiar, form. We do this to provide further visual contrast between the two groups and show the distribution of remediation capacity around the mean.

But Is It Enough?

We have established that remediation capacity can be expanded well above the norm, but we still do not know the all-important question of whether it's enough to actually win the race. To test that, we calculated the total number of open high-risk vulnerabilities and the number remediated per month for each organization. We then calculated a 'net remediation capacity' to determine which firms were keeping up (closing about as many vulnerabilities as were opened), falling behind (opened > closed), or gaining ground (closed > opened).

Figure 22 captures the results. From that, we see about half of firms falling behind, one in six keeping up, and one-third gaining ground. The figure also makes it plain that the degree by which organizations are falling behind or pulling ahead varies widely. This corroborates our findings above, suggesting that remediation capacity, though remarkably consistent across diverse organizational demographics, is not set in stone. It's possible to improve and exceed the norm. And yes—it is enough to win the remediation race.

FIGURE 22:

Comparison of net remediation capacity among firms.

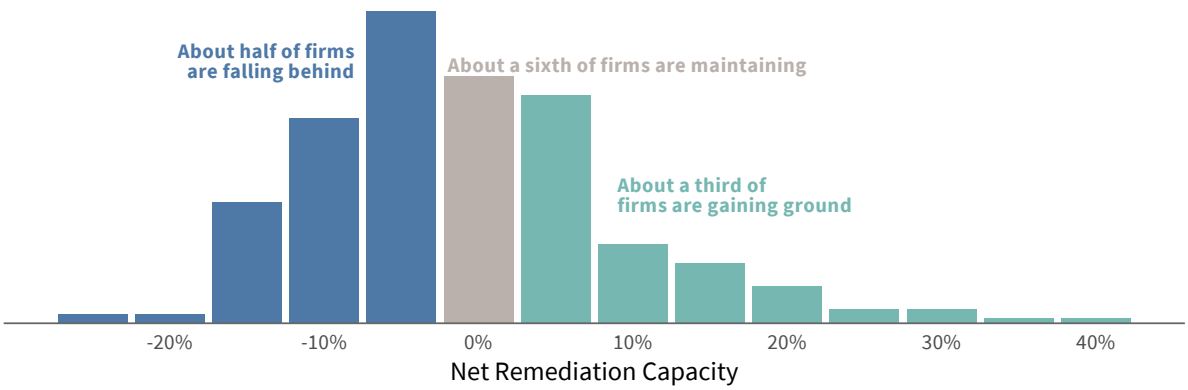

Source: Kenna / Cyentia

Reviewing What We've Learned

The three reports published thus far in the *Prioritization to Prediction* series span roughly 15,000 words, 60 pages, and 50 figures. Given all that content, one would hope that we've learned a thing or two about vulnerability exploitation and remediation. Indeed, we have!

We thought it would be fitting to close this third volume with a recap of major lessons from this research initiative thus far. To make it challenging, we set a goal to distill it all down to three question and answer pairs. Here goes:

Q Can organizations remediate new vulnerabilities before exploitation?

A: It depends on how you define exploitation. Volume 1 taught us that, if released, exploit code for 70% of (exploited) CVEs drops within the first month of publication. Compare that to Figure 3 above, which shows 25% of vulnerabilities are remediated in the first month, and it's clear that exploit writers have the edge on this one. Changing the definition to the start of exploitation in the wild buys some more time, but still doesn't erase that advantage. However, the chance that your organization is the first (or 1000th) target of exploitation is unlikely. Thus, we conclude that the majority of organizations have sufficient remediation velocity to address most vulnerabilities before they become victims of exploitation (assuming they know which to fix first).

Q Can organizations remediate all new vulnerabilities in their environment?

A: No; not by a long shot. Figure 19 above teaches us that organizations typically have the capacity to remediate one out of every 10 vulnerabilities in their environment. We're all in favor of improving that, but the simple fact of the matter is that the volume of new and existing vulnerabilities will always exceed the capacity to remediate them.

Q Can organizations remediate all new high-risk vulnerabilities in their environment?

A: Yes! Figure 22 proves that it is possible not only to keep up with, but to get ahead of the number of net new high-risk (exploited) vulnerabilities that pop up in your environment over time. Of course, that outcome depends on organizations having intelligence or guidance on which vulnerabilities are high-risk in order to prioritize them for remediation.

And there we go—three big lessons from three big reports. We hope these lessons, this new report, and this entire series help you to pull ahead in the relentless race of remediation. We'll be back with Volume 4 later this year. Until then, let us know how you're using this research and if there's anything we can do to help: @kennasecurity, @cyentiainst, #p2preport

Appendix A: Data Sources

This study focuses on the vulnerabilities described in MITRE's Common Vulnerabilities and Exposures (CVE) List. But in order to provide more context to the study and help measure the importance of remediating any specific CVE, we also leverage several other sources. We describe these sources and attributes in this section.

Common Vulnerabilities and Exposures (CVE)

We focus our research on discovered and disclosed vulnerabilities contained in the CVE List from MITRE. We do this primarily because CVEs are publicly tracked, readily available, extensive (although not exhaustive), and have become the de facto standard adopted by many other projects and products. It should be noted, however, that CVEs are neither comprehensive nor perfect. Many vulnerabilities are unknown, undisclosed, or otherwise have not been assigned a CVE ID. Furthermore, CVE listings are curated by humans, which makes them vulnerable to biases, errors, and omissions.[4] Despite these challenges, the CVE List is a valuable community resource that greatly assists the otherwise untenable task of vulnerability management.

Since its inception, well over 100,000 CVE entries have been created. Another 20,000+ are still in "reserved" status, meaning they have been allocated or reserved for use by a CNA or researcher, but the details have not yet been populated. Over 4,000 have been rejected for various reasons and another eight are split out or labeled as unverifiable. We chose to include the ~500 published CVEs currently in the "disputed" state since most describe weaknesses that would be useful to an attacker.

For all intents and purposes, each of these published CVEs represents a decision and potential action for vulnerability management programs. The criteria for those decisions may be simple in the singular case (e.g., "Does that exist in our environment?"), but prove to be quite difficult in the aggregate (e.g., "Where do we start?"). Figure 1 in our first report reinforces this challenge by demonstrating the increasing volume of reserved and published CVEs over time.

CVE Enrichment Projects

In addition to the basic CVE information produced by MITRE, this research also leverages the details added to each CVE by the National Vulnerability Database (NVD). NVD enriches the base CVE information with details leveraging other community projects, which include the following:

Common Vulnerability Scoring System (CVSS). This provides a process to capture the principal characteristics of a vulnerability and produce a numerical score that reflects its severity. The CVSS standard has very recently moved to version 3, but the majority of published CVEs were recorded using version 2, so we use version 2 in this report. CVSS was developed and is maintained by the Forum of Incident Response and Security Teams (FIRST).

Common Platform Enumeration (CPE). This provides a standard machine-readable format for encoding names of IT products, platforms, and vendors. It was developed at MITRE, but ongoing development and maintenance is now handled by NIST.

Common Weakness Enumeration (CWE). This provides a common language for describing software security weaknesses in architecture, design, or code. It was developed and is maintained by MITRE. We won't be discussing CWEs in this study.

4 *For discussion of these biases and other CVE-related issues, see 2013 BlackHat presentation titled "Buying into the Bias: Why Vulnerability Statistics Suck" from Brian Martin and Steve Christey.*

Each piece of enrichment data offers potentially useful context for decisions. Basic remediation strategies may rely on CVSS alone, while others will factor in the type of vulnerability (CWE) along with the vendor and product and the exposure of the vulnerabilities across environments.

Exploit Code and Activity

Basic analysis of CVEs may stop at data from MITRE and NVD, but those miss an important part of the equation: What CVEs are attackers actually exploiting in the wild? Unfortunately, no universal source of exploit activity exists, so we have to collect this information through multiple direct and indirect ways. These include host and network-based detection systems as well as by reverse engineering the malware and tools used by attackers.

Sources used in this study to track which CVEs have been actively exploited include the SANS Internet Storm Center (monthly statistics from the ISC signature collection Honeynet Project), Secureworks CTU (active campaigns associated with CVEs), AlienVault's OSSIM metadata (reputation feed, collecting IDS signature hits across 100,000+ devices in 150+ countries) and ReversingLabs metadata.

But sometimes observing exploitation in the wild comes too late for risk-averse vulnerability remediation strategies. In such cases, published exploit code serves as a good indicator of exploitability because it enables attackers to easily weaponize a vulnerability. Roughly two out of every three CVEs with active exploit detections also have published exploit code. Tracking the publication of exploit code, therefore, is important to remediation prioritization.

Sources used in this study to track which CVEs have public exploit code include Exploit DB, several exploitation frameworks (Metasploit, D2 Security's Elliot Kit, and Canvas Exploitation Framework), the Contagio dump and data from ReversingLabs, and Secureworks CTU.

Not only are exploit code releases strongly correlated with active exploitations, but they also indicate something more: the characteristics of a vulnerability that exploit writers target. Even if we haven't seen a specific CVE with published exploit code, the exploited vulnerabilities tend to share similar characteristics and traits with written exploits.

Vulnerability Observations

Information shared on observed vulnerabilities (open or closed) is drawn from the Kenna Security Platform. Hundreds of organizations use this platform as part of their vulnerability management programs. Ingested into the platform is a rich dataset from a variety of internal sources:

- Findings from any vulnerability scanner
- Asset and network-specific data from configuration management database (CMDB) tools
- Penetration test or red team findings
- Bug bounty programs
- Static application testing
- Dynamic application testing

Kenna uses all of this data to get a full view into the potential impact of each vulnerability, including the volume and velocity of attacker activity, as well as how critical each threat could be given your specific environment. We used it to derive the findings and statistics we shared in this report.

PRIORITIZATION TO PREDICTION

VOLUME 3: WINNING THE REMEDIATION RACE

KENNA Security

119 Cyentia INSTITUTE

> "It is possible not only to keep up with, but to get ahead of the number of net new high-risk vulnerabilities that pop up in your environment over time—if you have the right intelligence."