

PRIORITIZATION TO PREDICTION

Volume 8: Measuring and Minimizing Exploitability



This research was commissioned by Kenna Security, now part of Cisco. Kenna collected and provided the dataset to the Cyentia Institute for independent analysis for this report.

Kenna Security is the enterprise leader in risk-based vulnerability management. Kenna Security solutions enable organizations to work cross-functionally to determine and remediate cyber risks. They leverage machine learning and data science to track and predict real-world exploitations, empowering security teams to focus on what matters most. Headquartered in Santa Clara, Kenna serves nearly every major vertical and counts CVS, KPMG, and many Fortune 100 companies among its customers.

Find out more at www.kennasecurity.com.

Cisco Secure is built on the principle of better security, not more. It delivers a streamlined, customer-centric approach to security that ensures it's easy to deploy, manage, and use – and that it all works together. We help 100 percent of the Fortune 100 companies secure work – wherever it happens – with the broadest, most integrated platform. Learn more about how we simplify experiences, accelerate success, and protect futures at cisco.com/go/secure.

PRIORITIZATION TO PREDICTION

VOLUME 8: MEASURING AND MINIMIZING EXPLOITABILITY

Overview & Key Findings	2
Known Vulnerabilities and Exploits	3
Exploits of Disclosed Vulnerabilities	3
Exploits Targeting Vulnerable Assets	4
Measuring Exploitability	6
Exploitability of Published Vulnerabilities	7
Exploitability of Popular Products	8
A Survival Guide on Survival Curves	10
Exploitability within Assets	11
Minimizing Exploitability	14
Reviewing Remediation Capacity	14
Simulation Model Mechanics	15
Effect of Prioritization Strategies	17
Effect of Remediation Capacity	18
Why Not Strategy AND Capacity	19
Concluding Thoughts	20



Analysis for this report was provided by the Cyentia Institute. Cyentia is a research and data science firm working to advance cybersecurity knowledge and practice. We do this by partnering with security vendors and other organizations to publish a range of high-quality, data-driven content like this study.

Find out more at www.cyentia.com.

Overview & Key Findings

“It is essential to aggressively remediate known exploited vulnerabilities to protect federal information systems and reduce cyber incidents.”

—Cybersecurity and Infrastructure Security Agency, Binding Operational Directive 22-01

“If you take the vulnerabilities in your environment and focus on the ones that are being exploited in the wild, this will be an exponential improvement in your security posture.”

—Mitchell Schneider, Gartner, [Vulnerability Management – What is Working and What is Not](#)

It would be tough to set the stage for this eighth volume of Prioritization to Prediction (P2P) with statements more apropos than those quoted above from CISA and Gartner. Plus, the final paragraph of [Vol. 7](#) foreshadowed the plot for this report by asking, “Is it possible to determine the relative exploitability or remediability of an entire organization?” So we’ll skip the flowery introductory prose and just get straight to the point.

We do two very important and, based on the quotes above, timely things in this report. We first explore ways to measure exploitability for individual vulnerabilities—and far more importantly—entire organizations. Second, we create a simulation that seeks to minimize organizational exploitability under varying scenarios combining vulnerability prioritization strategies and remediation capacity. Bottom line: If you’re looking for proven ways to squeeze the most risk reduction from your vulnerability management (VM) efforts, this report is for you.

Key Findings



Nearly all (95%) assets have at least one highly exploitable vulnerability.



Twitter mentions offer a demonstrably better signal-to-noise ratio than CVSS does for remediating vulnerabilities most likely to be attacked in the wild.



Prioritizing vulnerabilities with exploit code publicly available is 11 times more effective than CVSS is for minimizing exploitability.



Staying focused is more important than fast fixes to efficiently reduce attack surface and minimize overall risk exposure.



Given the choice, it's far more effective to improve your strategy for prioritizing vulnerabilities than to increase remediation capacity.



Combining a good vulnerability prioritization strategy with high remediation capacity can achieve a 29X reduction in exploitability!

Known Vulnerabilities and Exploits

This section serves as a review and update for several things we've measured in prior P2Ps. It also sets up what's to come in this edition. When prioritizing remediation efforts, tracking disclosed vulnerabilities is an obvious starting point. As readers of this series are well aware, we also believe that tracking exploits targeting those vulnerabilities is essential. Together, they form the building blocks of risk-based vulnerability management (RBVM).

Exploits of disclosed vulnerabilities

Charting the number of new vulnerabilities disclosed to the [CVE List](#) or [National Vulnerability Database](#) (NVD) each year is table stakes for any statistical analysis on this topic. It drives home the points that 1) there are a lot of known vulnerabilities out there to track and 2) new ones are added at an increasing rate. Per Figure 2, we expect to see another ~18,000 vulnerabilities published to the CVE List for 2021—an average of 50 per day!

Moving beyond table stakes, Figure 1 conveys two important pieces of information that are NOT readily available from the CVE List. The first (teal) is the proportion of published CVEs observed in production enterprise assets, and the second concerns the ratio of those CVEs with known exploit code or active exploits in the wild (aka “high risk”). We'll address these in turn.




Figure 1: Number of vulnerabilities added to the CVE List annually with proportions that are observed and exploited.

We can all probably agree that tracking, evaluating, and remediating 18,000 new CVEs each year isn't even remotely feasible.[†] There's a lot of evidence from prior P2P reports to backup that statement, and you'll see some more of that in this one. The good news from Figure 1 is that we don't need to fret over them all because only about one-third of published CVEs are ever detected by a scanner in enterprise environments. And the proportion observed in YOUR environment is ostensibly much less than that. So Step 1 in reducing the vulnerability firehose is to filter the flow down to just the assets you're managing.

Unfortunately, the "asset filter" approach still leaves too many vulnerabilities to deal with based on the remediation capacity and velocity metrics we've established over time. We need another finer-grained filter. It turns out that focusing on high-risk vulnerabilities with known exploit code offers an excellent way to accomplish that. Our vulnerability intelligence identifies exploit code or activity for about 16% of all vulnerabilities on the CVE List.

But we can do even better than that in terms of focusing VM efforts on what really matters. If we apply the two filters described above for "observed" AND "high risk" CVEs, we're left with just over 4% of published vulnerabilities that represent a real risk to organizations (see the red portion in Figure 1). Not exactly a trickle, but much more manageable than the CVE List firehose we started with. Yay for intelligence and statistics!

We frequently refer to "high risk" vulnerabilities in this report.

This designation means we have intelligence that exploit code (i.e., a proof-of-concept or ready-to-use tool) exists or that the vulnerability has been exploited in the wild (i.e., used to probe and/or attack organizations).

Exploits targeting vulnerable assets

Reducing the number of disclosed CVEs to keep on the radar is helpful, but solves only part of the problem. If you've ever seen the Original Star Trek episode "[The Trouble with Tribbles](#)," you have a perfect on-screen analogy for vulnerabilities in a live environment. Individual CVEs can be interesting—even kind of cute when accompanied by a nifty logo—but they quickly multiply out of control when they infest assets across the *Enterprise*.[‡] Figure 2 puts some numbers and rule-of-thumb annotations around that statement.

[†] If you don't agree with this statement, please tell us your secret so we can productize it and get filthy rich together while saving the world.

[‡] We're claiming bonus points for the double entendre of "Enterprise" here. This is a tough job, folks; we have to Klingon to all the kudos we can get.




Figure 2: Distribution of the total number of assets affected by CVEs

If we roll the asset-centric tally of vulnerabilities from Figure 2 up to the organizational level, we arrive at the view in Figure 3. Here we see that most (87%) organizations have open vulnerabilities in at least a quarter of their active assets, and 41% of them show vulns in three of every four assets.

Even more telling, following the red “high risk” line reveals that 75% of firms have more than 1 in 4 assets that can be easily exploited. That risky stat jumps to at least 3 in 4 assets for 19% of organizations. So, while high-risk vulnerabilities might be rare among published CVEs, they are quite common among active assets.




Figure 3: Proportion of assets with open vulnerabilities among ~500 firms

Measuring Exploitability

Having intelligence on which vulnerabilities have been exploited is awesome. You can prioritize fixing those, thereby reducing exposure much more efficiently.[†] But what about vulnerabilities for which there are no known exploits or current attacks in the wild? Do we treat all of those the same? Assuming we're able to remediate vulnerabilities with known exploits and have excess capacity, which ones warrant attention next?

Being able to remediate all high-risk vulnerabilities may seem like a pipe dream, but you should know that quite a few organizations are living that dream. The majority (60%) of Kenna customers reduced the average number of open high-risk vulnerabilities in their environments over the last two years. Another 21% held their ground, meaning more than 7/10 firms are successfully managing vulnerability risk in the real world!




Figure 4: Net remediation capacity for high-risk vulnerabilities among firms

This is where the concept of “exploitability” comes into play. What's the likelihood that a given vulnerability will be exploited within a window of time? If that sounds like a hard question to answer, you're right. It is. But the good news is that Kenna Security, the Cyentia Institute, and quite a few other organizations have been collaborating on this exact thing. It's called the [Exploit Prediction Scoring System](#) (EPSS), and it's maintained by a Special Interest Group at FIRST.org.

[†] We compare common vulnerability remediation strategies in [P2P Vol. 1](#). If you'd like to see exactly how much better prioritizing exploited vulnerabilities works than, for instance, fixing all vulns with a CVSS score of 7 or above, check it out.

EPSS is an open, data-driven effort for predicting whether and when vulnerabilities will be exploited in the wild. The goal is to assist cybersecurity defenders to better prioritize remediation efforts. EPSS uses current information about CVEs combined with real-world exploit data from multiple sources to build a model that generates a probability between 0 and 1 (0 and 100%). The higher the score, the greater likelihood that the vulnerability in question will be exploited.

Let's see what EPSS scores can teach us about the exploitability of production assets across enterprise environments.

Exploitability of published vulnerabilities

Following the pattern of the previous section, we'll start by looking at EPSS scores across the published vulnerabilities on the CVE List. It's obvious from Figure 5 that the probability of any random disclosed vulnerability being exploited in the wild is quite low. A strong 63% majority of CVEs have less than a 1% chance of exploitation. Only 5% of CVEs exceed 10% probability. Of course, most of us aren't worried about any random vulnerability; we're worried about the ones that potentially affect us.




Figure 5: Distribution of EPSS scores among published CVEs

With that in mind, let's look at EPSS scores across vulnerabilities detected by vulnerability scanners in corporate environments. The "mound" of the distribution shifts noticeably to the right in Figure 6, indicating a higher likelihood of exploitation among vulnerabilities in software that organizations actually use. That's not terribly surprising, since exploit developers will generally want to target a large attack surface rather than a minuscule one.[†]




Figure 6: Distribution of EPSS scores among CVEs observed in production assets

[†] Yes, targeted attacks defy this generalization, but they still conform to the probabilities shown here.

The takeaway here is that the probability of exploitation for any given vulnerability disclosed publicly is pretty low. But EPSS scores rise significantly for vulnerabilities observed in live environments, suggesting a need to monitor the exploitability of assets over time to limit exposure.

Exploitability of popular products

Of course, not all assets are exploited equally. Software and hardware products differ widely based on install base, number of publicly disclosed vulnerabilities, exploit code development, etc. To illustrate that, Figure 7 compares exploitability among 80 common vendors.

The x-axis in Figure 7 corresponds to the prevalence of vulnerabilities associated with each vendor across all assets. On the vertical axis, we've averaged the EPSS scores of all published CVEs for each vendor.[†] Vendors toward the top are more likely to be the target of exploitation activity. Thus, vendors in the upper right quadrant (e.g., Microsoft, Adobe) represent a large and attractive attack surface across organizational assets.




Figure 7: Vulnerability prevalence and exploitability among common vendors

You can view Figure 7 as comparing inherent or baseline exploitability without regard to how firms are remediating those vulnerabilities. That may not capture the true state of things. Figure 8 addresses that by showing the rate of remediation[‡] on the x-axis.

[†] Yes, the distribution of scores is skewed (Figure 6), so the average wouldn't normally be the best measure of centrality. But in this case, the mean "penalizes" vendors with consistently higher scores, which seems appropriate for a comparison of relative exploitability.

[‡] The Area Under the Curve (AUC) is a measure of survival time in [survival analysis](#). A lower AUC is better because it means vulnerabilities are "dying off" (being remediated) more quickly.

This view paints a decidedly more favorable picture of Microsoft. It still tops the exploitability scale, but its position on the far left of the grid reminds us that Microsoft vulnerabilities are fixed more quickly than that of almost any other vendor. That's especially impressive given their huge install base.[†]

We could make many other observations from Figure 8, but we'll leave it to you to follow your own curiosity. See you on the other side.




Figure 8: Vulnerability prevalence and exploitability among common vendors

Also noteworthy is that Google is in a class of its own in terms of low exploitability and high remediation velocity.

“It’s clear certain vendors or products will make managing exploitability a more challenging feat. That’s not to say your security destiny is determined solely by the software in your environment. But Figures 7 & 8 definitely shows that VM programs which adapt to the strengths and weaknesses of their tech stack are best positioned to reduce risk efficiently.” (Adapted from P2P Vol. 7)

[†] You can get more details and reasons for Microsoft’s impressive track record on remediation in [P2P Vol. 5](#).

A Survival Guide on Survival Curves

Figure 8 presents Area Under the Curve (AUC) as a measure of remediation rates among vendors. You might be curious what a survival curve looks like, so we've created one that depicts how quickly organizations kill off all observed vulnerabilities across their assets. Overall, about half of vulnerabilities are remediated in the first month, but the long tail of remediation leaves 16% open for more than a year.




Figure 9: Overall remediation velocity across active assets for all firms and vulnerabilities

From the AUC comparisons back in Figure 8, you can probably surmise that survival curves for specific vendors or products would look very different from the one depicted above. No need to struggle in visualizing that; we've created mini survival curves below for five select vendors representing the spectrum from fast-to-slow remediation (AUC is in the upper right corner). Google's mobile- and browser-heavy portfolio achieves a remediation rate that's roughly 1/6 that of Schneider Electric's industrial control systems lineup which is likely embedded in the field and requires physical access to patch.




Figure 10: Comparison of vulnerability survival curves among five example vendors

Exploitability within assets

We'll now turn from measuring the exploitability of published CVEs to that of enterprise assets. This is where the rubber meets the road for organizations looking to reduce their attack surface. A concept presented in [prior P2P reports](#)—vulnerability density—is helpful here. Vulnerability density measures the number of vulnerabilities that exist within an asset and is inclusive of the base operating system, flaws in other installed software, and any other misconfigurations identified by a vulnerability scanner. To measure the exploitability of an asset, our calculation needs to consider all the vulnerabilities affecting it.

If you're thinking EPSS might help with that, good on you. You're paying attention. But EPSS calculates the probability of exploitation for a single vulnerability—not all that affect an asset. To account for that, we identified vulnerabilities with the highest EPSS score on each asset. This offers a reasonable way of measuring exploitability from an attacker-centric "just need one opening" perspective. Figure 11 presents this view of exploitability across all active assets in our dataset. And it's a fairly sobering view.




Figure 11: Distribution of maximum EPSS scores for vulnerabilities in all assets

Nearly all (95%) assets have at least one highly exploitable vulnerability with greater than a 10% chance of exploitation in the wild. Those may not seem like bad odds, but Figure 5 reminds us that an EPSS score of 0.1 (10% probability of exploitation) is in the top 5% of all disclosed CVEs. So, pretty much every asset in every organization has open vulnerabilities ranking in the upper echelon of exploitability.

If you take this finding to its logical conclusion, the next statistic and chart aren't really necessary. But we're going to share them anyway just to hammer home the point. Since EPSS represents a probability, we can combine the scores to calculate the likelihood that at least one open vulnerability within an asset will be exploited. Over half of all assets (1.8M of 3.5M) have a near-certain chance (>99%) that at least one of the open vulnerabilities will be actively attacked in the wild.[†]

[†] Keep in mind this is not the probability that a specific asset or organization will be exploited, just that the vulnerability will have observed exploitation activity somewhere across the internet.




Figure 12: Proportion of assets with at least one highly exploitable vulnerability.

Now that we can measure exploitability at the asset level, we're finally ready to return to the question that closed out [the last P2P report](#): *"Is it possible to determine the relative exploitability of an entire organization?"* It's easy enough to roll up the asset-level probabilities, but we hate to break it to you—the odds aren't in your favor. It's a near certainty that any organization with two or more assets will have a vulnerability in their environment that is (or soon will be) exploited somewhere in the wild. We won't bother with a chart for that one.

That being said, let's keep pulling on that thread. Is there perhaps a more useful way of measuring relative exploitability across an organization's attack surface? We think so and present that approach in Figure 13. Each dot represents an organization, and the horizontal axis tallies the number of active assets in the environment. Placement along the vertical axis is based on calculating the maximum EPSS score for each asset and then averaging those scores across the organization.

"Nearly all (95%) assets have at least one highly exploitable vulnerability. It's a near certainty that any organization with two or more assets will have a vulnerability in their environment that is (or soon will be) exploited somewhere in the wild."

We see two important takeaways from Figure 13. First, there's definitely wide variation in overall exploitability (or attack surface) among organizations. That's intuitive but helpful to see supported by the data. And it provides motivation to determine one's place in the pecking (or should we say "hacking") order. Second, exploitability doesn't appear to correlate with asset count. We don't think it's a stretch to conclude that firms of all sizes stand to benefit from better measuring and managing their attack surface.



Figure 13: Comparison of asset count and exploitability score for each organization

This is a good point for us all to take a breather and remember what's in sight here. Regardless of where it lands in Figure 13, we're not saying that **your** organization or its assets would or wouldn't be attacked. But we are saying that the attack surfaces of some organizations are larger and more exploitable than others. And it's possible to measure it for your organization. It's unlikely that your organization will be the very first target when a vulnerability starts getting exploited in the wild...but it probably won't be the last either.[†] Understanding the reality of exploitability across all your assets is crucial for prioritizing remediation efforts to minimize the likelihood of attacks succeeding whenever they do come your way.

“We see two important takeaways from Figure 13. First, there’s definitely wide variation in overall exploitability (or attack surface) among organizations. Second, exploitability doesn’t appear to correlate with asset count, so firms of all sizes stand to benefit from better measuring and managing their attack surface.”

[†] See *Timelines in the Vulnerability Lifecycle* in [P2P vol. 6](#).

Minimizing Exploitability

How can an organization most efficiently reduce exploitability across its attack surface? That's the question we take up in this final section, and we do so via a simulation model. The goal of the simulation is to minimize the exploitability score presented in the previous section using different remediation strategies.

To make that more challenging, we introduce remediation capacity as a constraint on how many vulnerabilities can be closed in a given timeframe. When all is said and done, we'll be able to answer questions like *Should I increase capacity or change my priorities? Which strategy works best? Should I listen to Twitter or CVSS?* Inquiring minds want to know, so let's get started.

Reviewing remediation capacity

Over the entire span of the [P2P series](#), we've repeatedly stated and shown that organizations cannot fix all the vulnerabilities across all their assets all of the time. But we've also demonstrated that some have more capacity to remediate than others. Remediation capacity measures the average proportion of open vulnerabilities closed in a given time period. It's a concept we first introduced in [P2P Vol. 3](#) with a chart just like the one below.




Figure 14: Average number of observed vs. closed vulnerabilities per month per firm

Every dot in Figure 14 marks the number of observed and closed vulnerabilities for a particular organization. Following the trendline reveals something close to a 1-in-10 fixed-to-observed ratio across the grid (it's actually 15%, which is an improvement over Vol. 3). Some firms manage to beat those odds (above

the line) and others fare worse, but that overall ratio is remarkably consistent across SMBs and large enterprises alike.

That's a really important finding, but the format of Figure 14 makes it hard to pinpoint the proportion of vulnerabilities closed per month. Figure 15 is better suited to that and presents the distribution of remediation capacity among organizations. The tail skews far out to the right (toward a smaller number of firms with very high remediation capacity), but the annotated points mark those in the middle of the road.




Figure 15: Distribution of vulnerability remediation capacity among organizations

The median remediation capacity across all organizations is 15% of open vulnerabilities in a given month. The lower quartile of firms fixes less than 6.6% of issues, while the upper quartile is able to exceed 27%. We'll use these low (25th percentile), mid (median), and high (75th percentile) breakpoints for the remediation capacity constraint in our simulation model.

Simulation model mechanics

As stated in the intro to this section, the goal of this simulation is to minimize exploitability using various prioritization strategies, subject to the organization's remediation capacity. Way back in the [OG of the P2P series](#), we compared the performance of several remediation strategies based on the level of coverage (recall) and efficiency (precision) they achieved. We're bringing some of those strategies into this simulation and adding some new ones:

Random: Randomly select vulns to fix. While a decidedly uneducated approach, it represents a baseline we always want to do better than.

CVSS: Prioritize vulns with the highest CVSS scores

Exploit code: Prioritize vulns whose exploit code is available

Prevalence: Prioritize the most observed vulns across all assets

Quickest: Prioritize vulns with the fastest remediation rates

Twitter: Prioritize vulns with the most mentions on Twitter

Perfect info: Prioritize vulns with the highest EPSS scores or known exploits in the wild. We'll use this as a proxy for having perfect information about what will be exploited.

To establish a starting point for organizations with respect to their current exploitability score and remediation capacity, we offer Figure 16. Each organization's placement along the x-axis corresponds to their calculated remediation capacity (see Figure 14). Our measurement of exploitability on the y-axis takes the maximum intelligence-modified EPSS score[†] for each asset and averages those scores across the organization (see Figure 13).

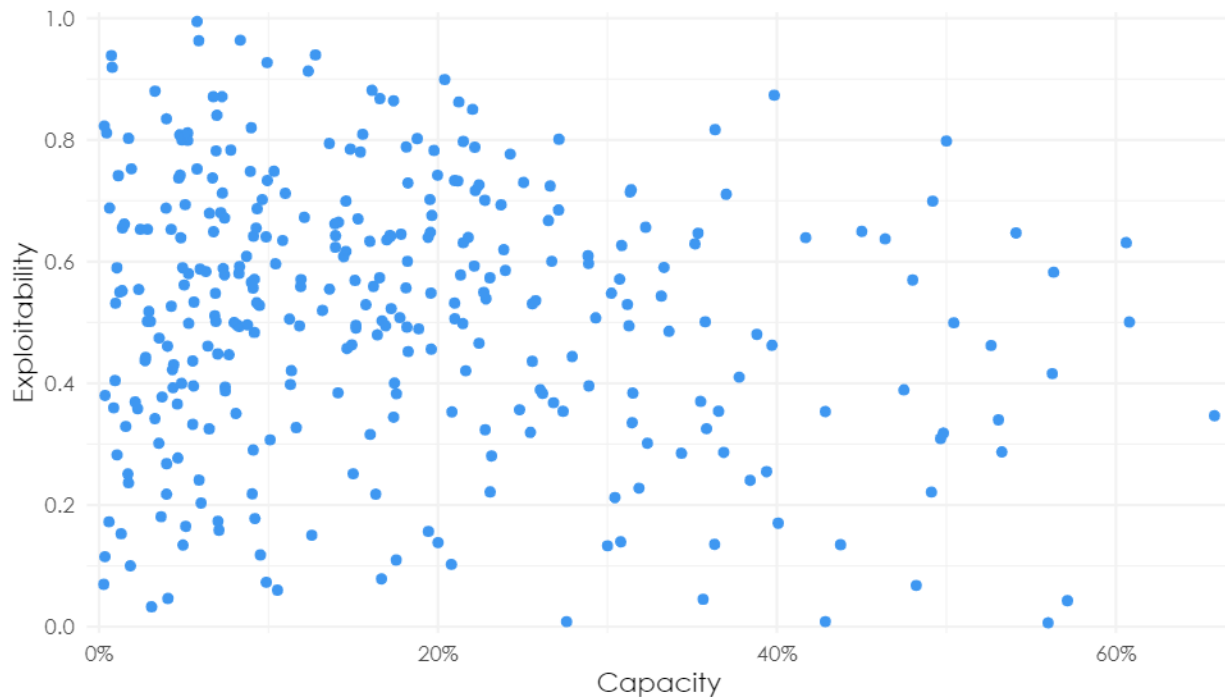


Figure 16: Comparison of remediation capacity and exploitability for each organization

In general, there doesn't appear to be a strong correlation between these two parameters. This suggests that minimizing exploitability isn't just a matter of doing more or doing it faster. We suspect it's also crucial to prioritize better, but we'll leave that to the simulation to confirm or deny. For now, let's simply observe that some low-capacity firms manage to achieve low exploitability, some with high remediation capacity are still relatively exploitable, and others fall everywhere in between.

In order to do a "what if" simulation across different prioritization strategies, we start with all of the currently open vulnerabilities at each organization and calculate the existing exploitability score. Then we apply the prioritization strategy at three different levels of capacity (low, median, and high from Figure 15).

Given the strategy and capacity, we can simulate the closing of prioritized vulnerabilities and recalculate the post-remediation exploitability score for each organization. For example, using a random strategy with low capacity, the simulation will randomly select and close 6.6% of discovered vulnerabilities. With a CVSS strategy, it closes vulnerabilities with the highest CVSS scores until the capacity is reached. This method is repeated across all of the strategies and capacity levels for all organizations to produce the results we'll discuss in the next sections.

[†] EPSS is a probability-based prediction, not a continuously-updated record of exploitation in the wild. For this simulation, we replaced the native EPSS score with 1.0 (certain probability) if we had intelligence that the vulnerability has been exploited in the wild.

Effect of prioritization strategies

To compare the effect of prioritization strategies on exploitability, we'll hold the remediation capacity constant at the median level. That means each organization can close 15% of its vulnerabilities in a given month. Figure 17 depicts the resulting exploitability scores for each organization (represented by the dots) under each prioritization strategy. The yellow dot marks the median exploitability score across all organizations.

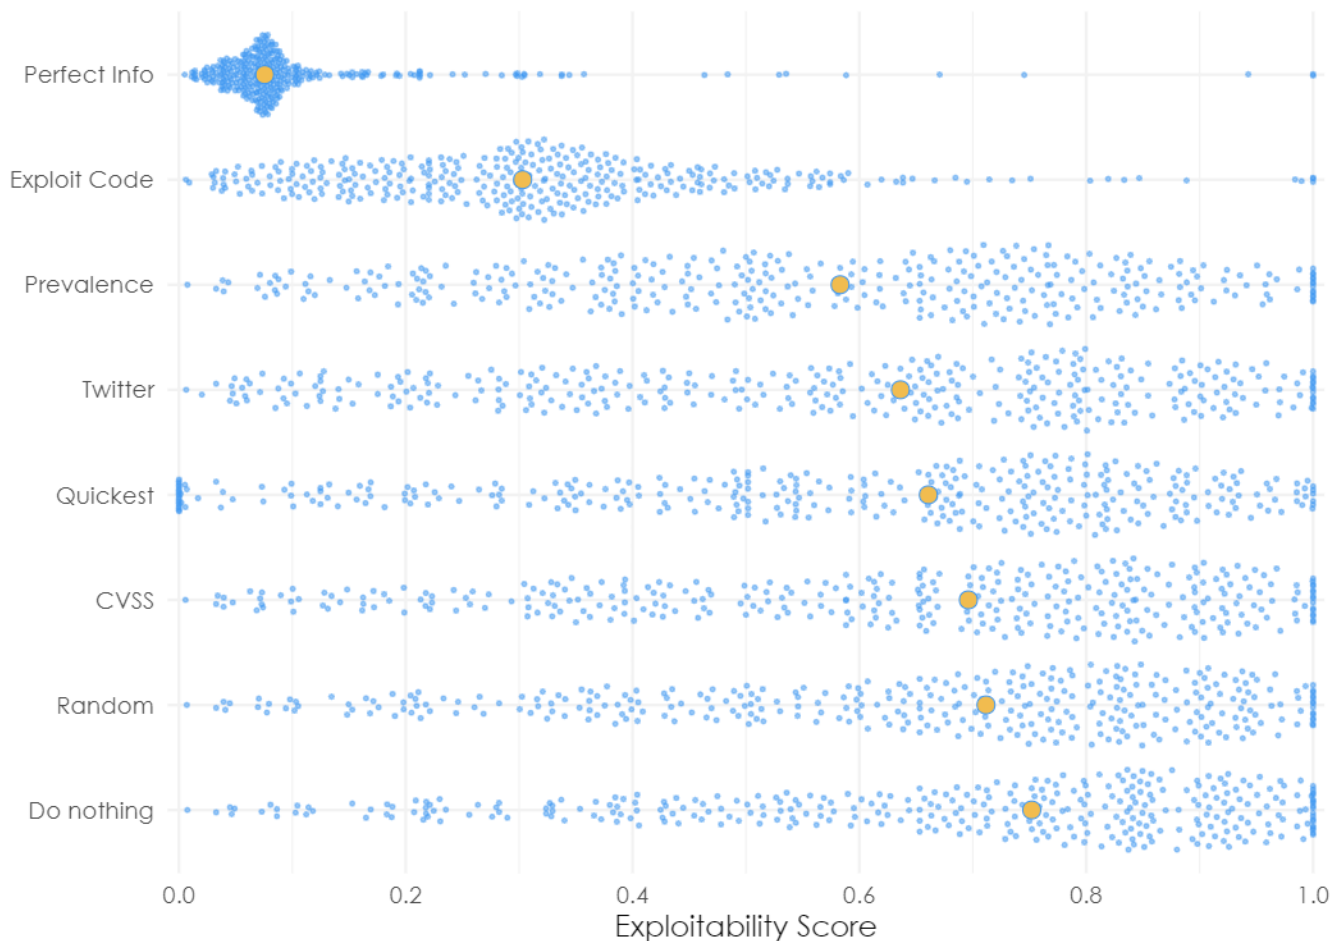


Figure 17: Effect of vulnerability prioritization strategies on organizational exploitability

Starting at the bottom, the Do Nothing strategy shows the fate of organizations that don't remediate any of the open vulnerabilities detected in their environments. Essentially, it coincides with the y-axis from Figure 16. That's obviously not a realistic remediation strategy, but it sets a baseline for improvements attributed to the other strategies. Jumping to the top, the Perfect Info strategy represents the theoretical best outcome if an organization could discern precisely which vulnerabilities would be exploited in the wild and focus exclusively on fixing those. Note that this still doesn't drive exploitability to zero because of capacity limitations.

Between those extremes, each prioritization strategy gets a little bit better than the one below it. We suspect some will be surprised to learn that prioritizing remediation based on CVSS scores performs about the same as randomly selecting vulnerabilities to fix. Perhaps even more surprising is that queuing off the raw count of Twitter mentions reduces exploitability more effectively than a CVSS-informed strategy.

“Perhaps even more surprising is that queuing off the raw count of Twitter mentions reduces exploitability more effectively than a CVSS-informed strategy.”

Targeting vulnerabilities with published exploit code proves to be an exceptional approach here. And if you've been following the [P2P series](#), that should come as no surprise. One of our key findings from Volume 7 was that vulnerabilities with exploit code rack up nearly 15X the overall exploitation activity in the wild compared to those without exploit code. This strategy requires reliable intelligence on exploit development but offers a strong signal-to-noise ratio for VM programs seeking to efficiently strengthen their attack surface.

Effect of remediation capacity

Now we'll examine how altering the remediation capacity affects simulation results. To do that, we calculate the median exploitability score achieved across all organizations at varying capacity levels for each strategy. Figure 18 presents the results for low (fix 6.6% of vulns/month), median (15.2%), and high (27.1%) remediation capacities.

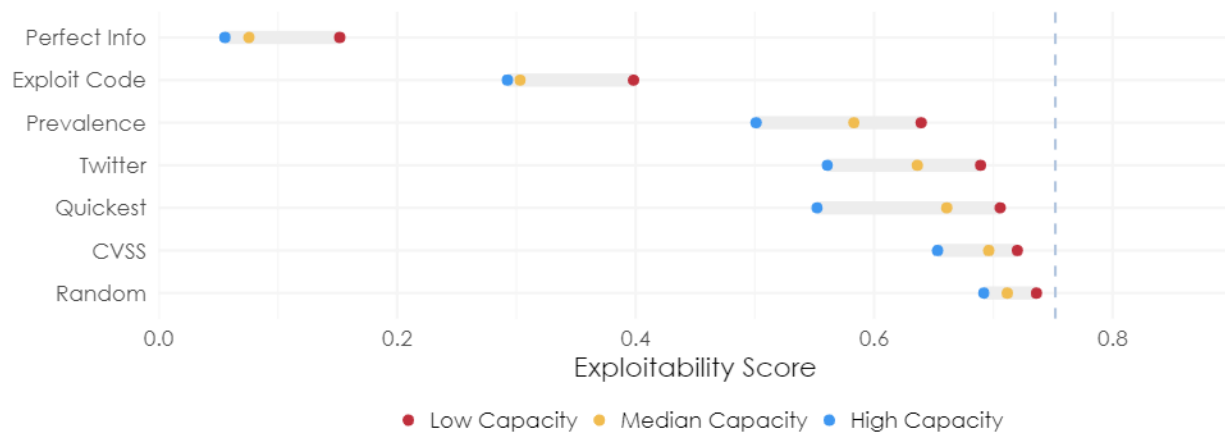


Figure 18: Effect of vulnerability remediation capacity and prioritization strategy on exploitability

The reduction in exploitability between low and high capacities differs among the strategies, but all show improvement. A key observation here is that some strategies ALWAYS outperform others, regardless of remediation capacity. In other words, prioritizing vulnerabilities with known exploits will reduce your organization's attack surface more effectively than quadrupling capacity to patch what CVSS deems to be critical. We thus infer that prioritization is more critical for managing risk than raw capacity is.

“Some strategies ALWAYS outperform others, regardless of remediation capacity. Prioritizing vulnerabilities with known exploits will reduce your organization's attack surface more effectively than quadrupling capacity to patch what CVSS deems to be critical.”

Why not strategy AND capacity?

Not everything has to be this OR that; sometimes you actually can have your cake and eat it too. We've shown that a good prioritization strategy can have a greater impact on exploitability than boosting remediation capacity can, but why not improve both? Let's do one more chart to demonstrate how you can make $1 + 1 = 29$ for your VM program.

Figure 19 lays out a compelling path to minimize exploitability across your enterprise. Compared to randomly fixing vulnerabilities as they pop up, low capacity with a poor strategy (CVSS) cuts your vulnerable attack surface in half. Sticking with CVSS and increasing remediation capacity from low to high triples that reduction to 6X. Alternatively, keeping capacity low while switching to a better strategy of prioritizing exploited vulnerabilities drops exploitability by 22X. So, if you have to make a choice, better prioritization can leapfrog capacity constraints to minimize risk for your organization.

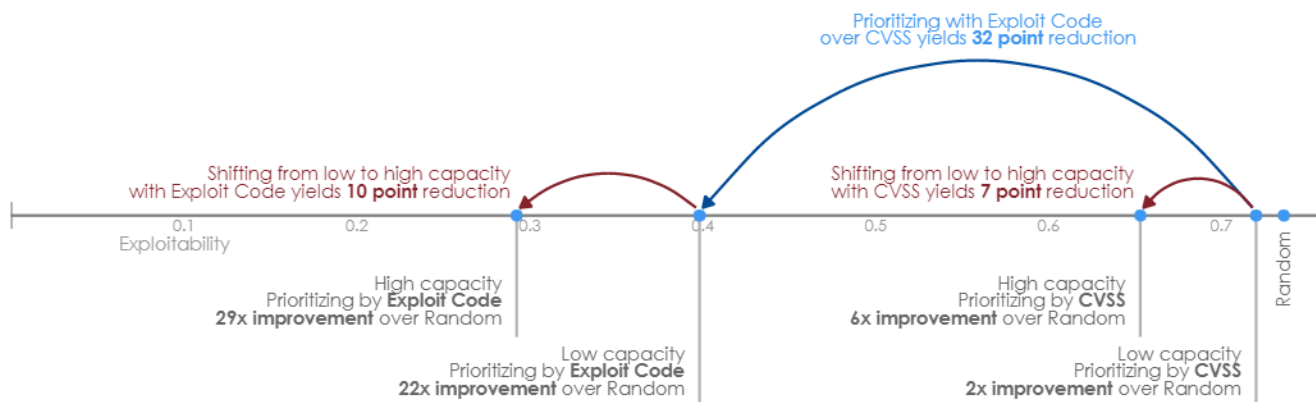


Figure 19: Exploitability reduction achieved by improving remediation strategy and capacity

But again, we ask, "Why not both?" An organization combining a good vulnerability prioritization strategy (Exploit Code) with high remediation capacity can achieve a 29X reduction in exploitability! And you thought we were just making up math when we said $1 + 1 = 29$ earlier. You should know us better than that by now.

Multiples like those shown in Figure 19 serve to hammer home the message that vulnerability management is not a mindless, endless loop of finding and fixing. Organizations have a great deal of control over their attack surface through the strategies and capabilities they employ. It's our hope that this latest volume in the P2P series puts some of that control back in your hands and helps you chart a course toward minimizing exploitability in your organization.

"Organizations have a great deal of control over their attack surface. Those combining a good vulnerability prioritization strategy (Exploit Code) with high remediation capacity can achieve a 29X reduction in exploitability!"

Concluding Thoughts

“Attackers do not rely only on ‘critical’ vulnerabilities to achieve their goals; some of the most widespread and devastating attacks have included multiple vulnerabilities rated ‘high’, ‘medium’, or even ‘low’. Known Exploited Vulnerabilities should be the top priority for remediation.” -CISA, *Binding Operational Directive 22-01*

“Vulnerabilities are nothing more than vectors for adversaries to control execution on a system...understanding how an adversary can, or is, leveraging this vulnerability in the wild can help with prioritization.” -*The Forrester Wave™ Vulnerability Risk Management, Q4, 2019*

And thus, this eighth volume of Prioritization to Prediction ends just like it began, with quotes from two influential institutions calling for organizations to evolve their approach to vulnerability remediation. We're delighted their statements echo the consistent finding from our research that prioritizing exploited vulnerabilities offers the most efficient and effective strategy for risk-based vulnerability management (RBVM).

Now that “Prioritization” is (or soon will be) the new norm, we'd like to raise the bar a bit further to encourage continued evolution of RBVM. The other “P” in the title of this series, “Prediction,” represents our stance that remediating known exploited vulnerabilities is a great start...but it's not enough. To truly get ahead of attackers and minimize exploitability, we need to accurately predict which vulnerabilities are most likely to be exploited in the future.

The good news is that we're well down that path through this P2P research series, community efforts like EPSS, and our own work on the Kenna Security platform. This volume demonstrates there's A LOT to be gained—or perhaps we should say a lot of risk to be lost—from prioritizing exploited vulnerabilities and predicting what attackers will target next. Thank you for joining us on this journey of discovery and development!

CISA recommends the following in Binding Operational Directive 22-01:

- Establish a process for ongoing remediation of exploited vulnerabilities.
- Remediate vulnerabilities according to the timelines established by CISA.
- Report on the status of exploited vulnerabilities in accordance with Continuous Diagnostics and Mitigation (CDM) requirements.

PRIORITIZATION TO PREDICTION VOLUME 8: MEASURING AND MINIMIZING EXPLOITABILITY

