tenable

PREDICTIVE PRIORITIZATION: DATA SCIENCE LETS YOU FOCUS ON THE 3% OF VULNERABILITIES LIKELY TO BE EXPLOITED

Contents

Executive Summary

In a perfect world, cybersecurity and IT professionals would proactively identify and patch every potential vulnerability, so their organizations could be protected from all known vectors of attack. But while digital transformation has given rise to new growth opportunities, it has also introduced uncharted areas of risk. Cloud, DevOps, mobile, IoT and critical infrastructure are now all aspects of the corporate cyberattack surface. And, traditional vulnerability management methods are simply no match for this dynamic environment. Nevertheless, cybersecurity and IT teams are under constant pressure to keep the organization safe from the vulnerability onslaught.

In 2017, 15,038 new Common Vulnerability and Exposures (CVEs) were published, up from 9,837 in 2016 – an alarming 53% increase.[1] In 2018, there were 16,500 new CVEs.[2] Since the average enterprise finds 870 CVEs across 960 IT assets every single day,[3] patching all vulnerabilities just isn't practical. Organizations must reduce the problem set to a manageable size. Many organizations use the Common Vulnerability Scoring System (CVSS) to rank order what should be patched. But, CVSS alone is limited in its ability to aid operational effectiveness. Instead, businesses need to know the difference between vulnerabilities representing theoretical versus actual risk – and then prioritize those vulnerabilities according to the risk level they pose.

> **This technical whitepaper explains the challenges cybersecurity professionals face, how they're prioritizing vulnerabilities today and how they can dramatically improve cyber risk management with Predictive Prioritization – the process of re-prioritizing vulnerabilities based on the probability that they will be leveraged in an attack.**

Predictive Prioritization combines more than 150 data sources, including Tenable® vulnerability data and third-party vulnerability and threat data, leveraging a proprietary machine learning algorithm to identify the vulnerabilities with the highest likelihood of exploitability in the near-term future. With Predictive Prioritization, organizations can dramatically improve their remediation efficiency and effectiveness by focusing on the 3% of vulnerabilities that have been or will likely be exploited.

16,500
CVEs in 2018

870
CVEs across 960 IT assets daily

1 CVE is maintained by the MITRE Corporation

2 National Vulnerability Database (NVD)

3 Tenable Research

Risk Management Challenges Today's Cybersecurity Professionals Face

According to Gartner, "Through 2021, the single most impactful enterprise activity to improve security will be mitigating vulnerabilities."[4]

But, mitigating vulnerabilities is harder than ever. The cyberattack surface has expanded beyond traditional IT-managed assets to include cloud, DevOps, mobile and web infrastructure – not to mention newly connected equipment like IP-enabled operational technology and IoT devices (see Figure 1). This growth in assets exacerbates the vulnerability overload problem.

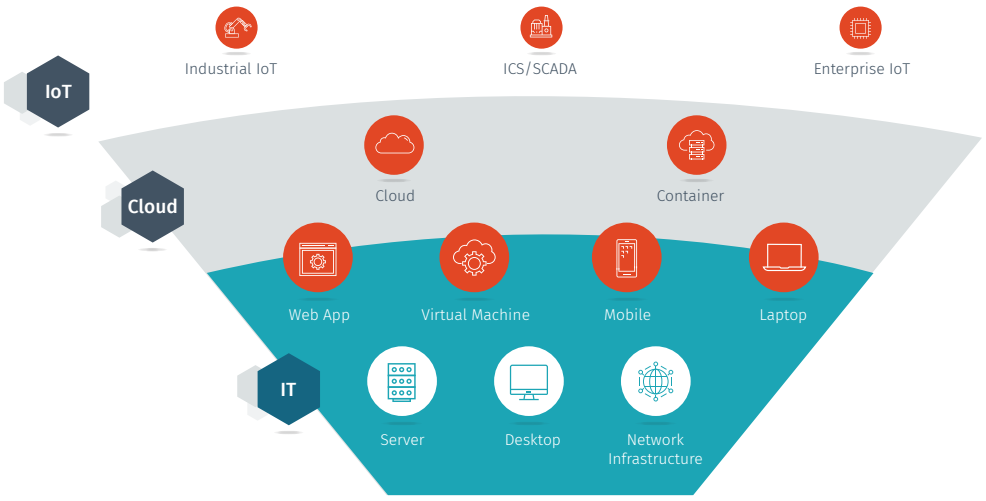

Figure 1. Digital transformation has multiplied the vectors of cyberattack

Gaining continuous, holistic visibility into this ever-shifting attack surface – and its many ephemeral assets (e.g., software containers) – is a fundamental problem in cybersecurity. But, once security teams solve that challenge and create a comprehensive view of all vulnerabilities, another conundrum follows: How to prioritize the thousands of vulnerabilities that superficially all seem the same?

4 Gartner Security and Risk Management Summit 2018 Presentation, "Fix What Matters: Provide DevOps Teams With Risk-Prioritized Vulnerability Guidance," Dale Gardner, June 4-7, 2018

Prioritization Is a Critical Issue Every Organization Faces

The following are prioritization-related problems that security teams routinely confront:

- **Vulnerability overload:** The discovery and disclosure of vulnerabilities continue to grow in volume and pace.
- **Triage based on High and Critical severity is futile:** The reality is, for most vulnerabilities, a working exploit is never developed. And, an even smaller subset of vulnerabilities are actively weaponized and employed by threat actors.
- **Distractions from real risks:** High-profile and zero-day vulnerabilities are often perceived as bigger threats than the risk they actually represent.

> **"[With] any large network, I will tell you that persistence and focus will get you in, will achieve that exploitation without the zero days. There's so many more vectors that are easier, less risky and quite often more productive than going down that route."**
>
> **– Rob Joyce, Senior Advisor for Cybersecurity Strategy to the Director of the National Security Agency (NSA)**

DID YOU KNOW?

Findings from Tenable Research's *Vulnerability Intelligence Report*

Vulnerability overload

15,038[5] new vulnerabilities were published in 2017 versus 9,837 in 2016 – a 53% increase in a single year. In 2018, that number jumped to approximately 16,500 more vulnerabilities.[6]

Triage based on High and Critical severity is futile

Prioritization methodologies based on remediating only Critical CVEs still leave the average enterprise with 100+ vulnerabilities per day to prioritize per patch.

5 Vulnerability Intelligence Report, Tenable Research, 2018

6 National Vulnerability Database (NVD)

Why Common Approaches to Vulnerability Prioritization Are Ineffective

The Pitfalls of Using CVSS to Prioritize

Most organizations use CVSS to prioritize vulnerability management efforts. It is a free and open industry-standard framework for assessing the severity of vulnerabilities. Organizations use CVSS to prioritize responses and resources based on numerical values from 0 to 10, where 10 represents the highest level of criticality.

However, the shift from CVSSv2 to CVSSv3 has had a huge impact on the distribution of severity:

- CVSSv3 scores the majority of vulnerabilities as High (7.0–8.9) and Critical (9.0–10.0).
- CVSSv2 scores 31% of CVEs as High severity (7.0–10.0), versus 60% with High (7.0–8.9) or Critical severity (9.0–10.0) under CVSSv3.[7]

Figure 2 (see below) illustrates the impact of this reclassification, underscoring a clear decline in the number of Medium CVEs, a spike in the number of High CVEs and a significant number of Critical CVEs (a new severity rating under CVSSv3).

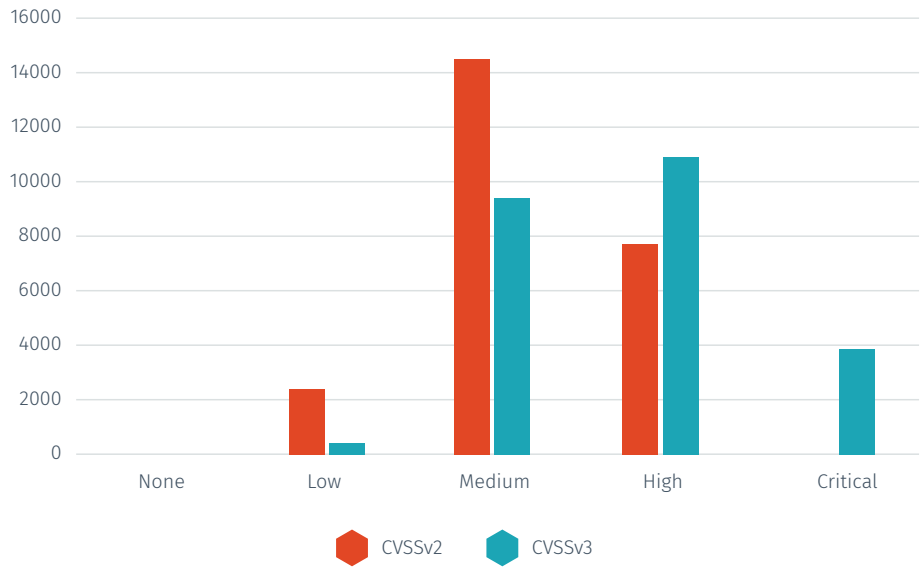

Figure 2. CVEs overall - CVSSv2 vs CVSSv3 classification

7 Vulnerability Intelligence Report, Tenable Research, 2018

In summary, CVSS falls short as a prioritization metric because it lacks sufficient granularity to distinguish between the highest degrees of severity – if everything is "critical," how can anything be?

CVSS + Exploit Database: The Rearview Mirror Approach

Since CVSS scores (particularly CVSSv3 scores) result in a body of vulnerabilities too large to handle effectively, another common strategy for prioritization is to target vulnerabilities with publicly available exploit code. Augmenting CVSS data with vulnerabilities appearing on Exploit Database, for example, will greatly reduce the quantity of vulnerabilities that need to be prioritized for remediation. For example, in 2018, only 7% of the 16,500 vulnerabilities had a public exploit available (see Figure 3).[8]

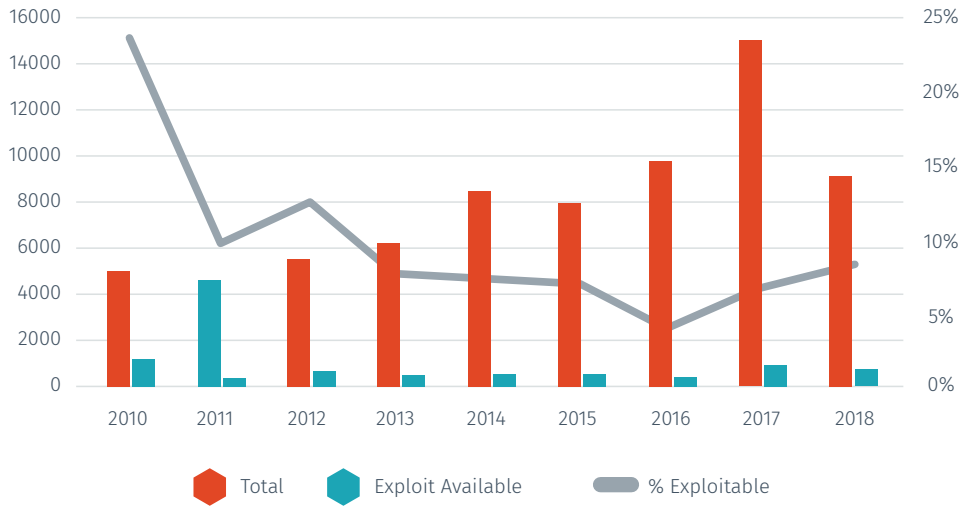

Figure 3. Total CVEs vs. exploitable CVEs

Note: At the time of the report's publication, 2018 data was not complete; this graph is based on disclosed CVEs through August 2018.

Exploit Database and Metasploit are useful sources of information regarding known and active threats. However, they don't provide insight into which vulnerabilities are likely to be exploited in the near-term future. Predictive processes are needed based on the characteristics of a given vulnerability.

8 Vulnerability Intelligence Report, Tenable Research, 2018

Introducing Predictive Prioritization: Turning Vulnerability Data into Actionable Intelligence

Predictive Prioritization is the process of re-prioritizing vulnerabilities based on the probability that they will be leveraged in an attack. Predictive Prioritization combines over 150 data sources, including Tenable vulnerability data and third-party vulnerability and threat data, leveraging a proprietary machine learning algorithm to identify the vulnerabilities with the highest likelihood of exploitability in the near-term future. With Predictive Prioritization, organizations can dramatically improve their remediation efficiency and effectiveness by focusing on the 3% of vulnerabilities that have been or will likely be exploited.

The Output of Predictive Prioritization: Vulnerability Priority Rating (VPR)

The output of the Predictive Prioritization process is called the vulnerability priority rating (VPR). VPR is a number that indicates the remediation priority (0–10, with 10 being the highest severity) of an individual vulnerability. For example, a vulnerability currently being exploited on a widely deployed service would have a significantly higher VPR than a vulnerability where no working exploit has been observed.

VPR is a dynamic number and changes with the threat landscape. It enables organizations to focus on remediating the vulnerabilities with the highest likelihood of being leveraged in a cyberattack. There are three components that make up the VPR:

1. **CVSS impact score**
2. **CVSS scope**
3. **A machine learning-based threat score**

Predictive Prioritization Makes Smarter Use of the CVSS Framework

CVSS does a good job capturing the scope and impact of vulnerabilities; it offers a sound explanation of what could happen if a given vulnerability is exploited. It also provides a foundation to gauge the likelihood of a vulnerability being exploited. However, its current application fails to deliver the granularity needed to prioritize effectively.

Predictive Prioritization remains true to the CVSS framework (see Figure 4), but enhances it by replacing the CVSS exploitability and threat components with a threat score produced by machine learning – powered by a diverse set of data sources. This means organizations can make remediation decisions based on the vulnerabilities that:

- Are likely to be exploited
- If exploited, will have a major impact

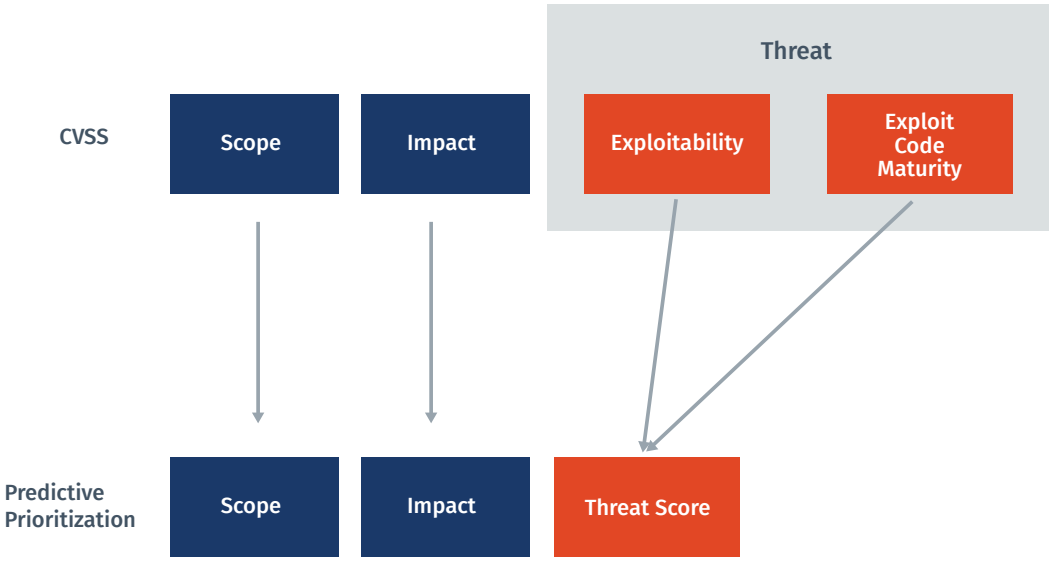

Figure 4. CVSS to Predictive Prioritization Framework

Note: Because Predictive Prioritization is based on the CVSS framework, organizations can benefit from Predictive Prioritization while using any CVSS-reliant processes they may have in place.

How Does Predictive Prioritization Work?

The Predictive Prioritization model takes more than 150 different aspects of the vulnerability into account. These characteristics are grouped into seven categories (see Figure 5):

- Past threat patterns
- Past threat sources
- Vulnerability metrics
- Vulnerability metadata
- Past hostility
- Affected vendor
- Exploit availability using threat intelligence data

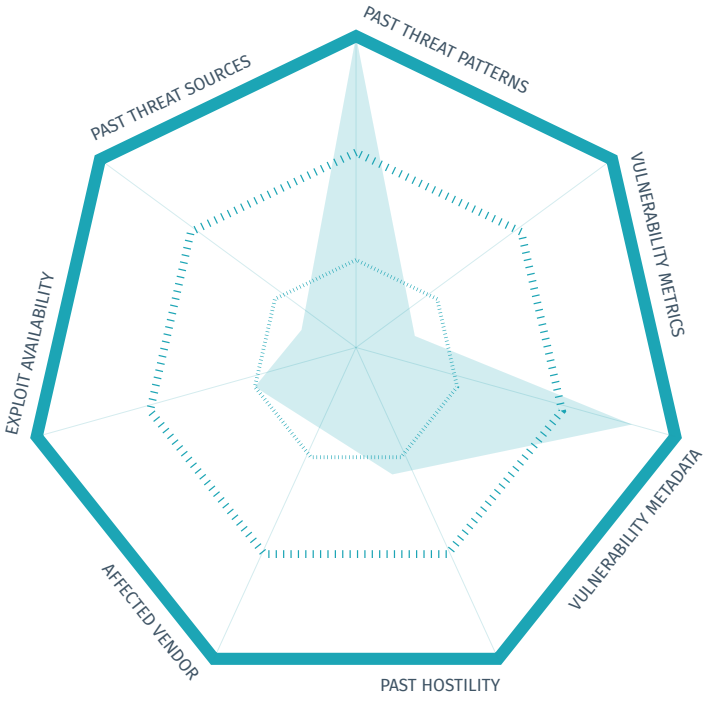

Figure 5. Predictive Prioritization

150

Different aspects to the model, in 7 different categories.

100K+

Priority calculated nightly on over 100,000 different vulnerabilities being tracked.

97%

Fewer vulnerabilities requiring remediation with the same reduction in attack surface.

After analyzing these 150 characteristics, Predictive Prioritization then scores more than 100,000 vulnerabilities being tracked[9] to identify the vulnerabilities with the highest likelihood of exploitability in the near-term future.

Threat Model Development

Predictive Prioritization gathers data across an ever-growing list of threat intelligence sources. It then uses natural language processing and other feature engineering techniques to extract the valuable components from this vast sea of data. Finally, multiple predictive, machine-learning models work together to produce the threat score (see Figure 6).

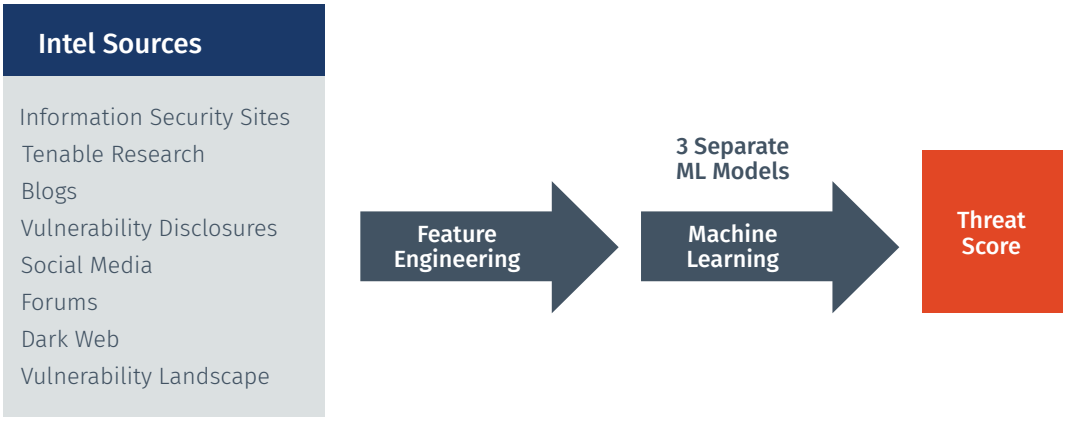

Figure 6. Tenable Threat Model Development

9 National Vulnerability Database (NVD)

An automated process analyzes the extensive amount of data on a given vulnerability and generates a list of features related to that vulnerability – including its age, availability of exploit kits, presence in Exploit Database, whether it's a topic of discussion on the dark web, forums and/or social media, etc. The model balances and weighs all these factors and constantly reassesses these metrics to reflect changes in the characteristics of the vulnerability.

Here are a few key principles behind how the model works:

- It scores all vulnerabilities with a CVE, updating the scores daily
- It is continuously monitored and retrained as needed
- Additional data is incorporated into the modeling process as it becomes available

The model uses historical data to understand the relationship between the input features and the likelihood of threat. This relationship is learned automatically during model training, and avoids the need to create and maintain long lists of complex rules that may be biased and flawed from the start.

With this dynamic model, vulnerabilities can be scored daily, which means the score on any given day represents the real-time threat risk as the threat landscape changes and evolves. Additionally, the model builds on the existing CVSS framework to produce a single score that combines threat intelligence and exploit code maturity, providing a complete view of the threat.

Predictive Prioritization is extremely responsive, giving you fast insights into the vulnerabilities that represent a real threat. See example (CVE-2017-0199: HTA Handler Vulnerability):

	Impact	Exploitability	Total
CVSSv3	5.9	1.8	7.8

	Impact	Threat	Total
Predictive Prioritization	5.9	3.9	9.8

Predictive Prioritization's Dynamic Scoring Offers Key Advantage Over Static CVSS

Another problem with using CVSS scores alone is a vulnerability's score may not change over time, even though the criticality of the CVE may vary as exploit kits are written and attacks are made on that vulnerability.

The actual criticality of a vulnerability may shift with related activity, but that does not mean the vendor responsible for reporting the vulnerability will update the CVSS score. And, even if they do revise the score, it still may not be accurate since circumstances can change rapidly and often.

As an example, the chart below shows the history of the Linux Kernel vulnerability known as CVE-2018-17182 (see Figure 7). Notice how its CVSS score remains unchanged while its Predictive Prioritization score adapts to associated events.

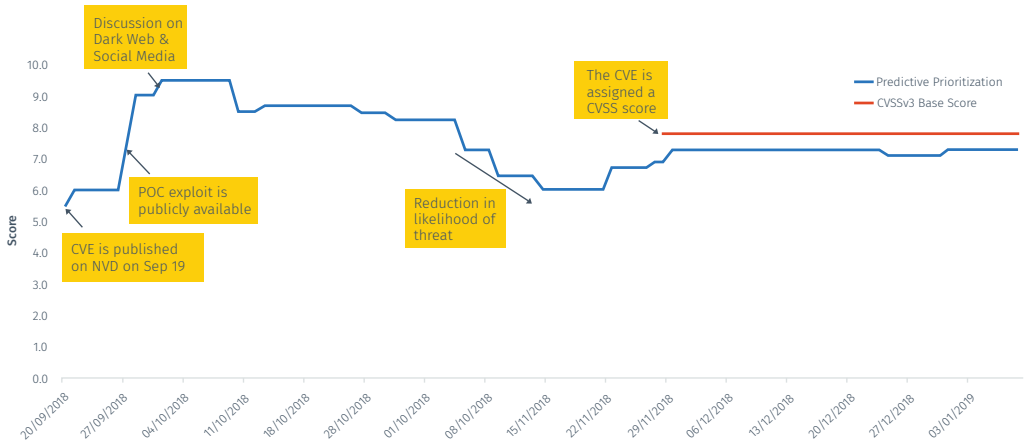

Figure 7. Dynamic Nature of Predictive Prioritization: CVE-2018-17182

Chart Timeline

- September 19, 2018: CVE published to NVD
- September 26, 2018: Exploit Database entry created
- September 30, 2018: Dark web and social media discussions
- November 5, 2018: Likelihood of threat starts to reduce
- November 28, 2018: CVE assigned CVSS vector/scores

Chart Discussion

The CVE is published on NVD on September 19, 2018, but has not yet been assigned a CVSS vector (and consequently a score). Predictive Prioritization uses a machine learning model to predict the CVSSv3 impact score based on the raw text found in the NVD description of the vulnerability:

> **"An issue was discovered in the Linux kernel through 4.18.8. The vmacache_flush_all function in mm/vmacache.c mishandles sequence number overflows. An attacker can trigger a use-after-free (and possibly gain privileges) via certain thread creation, map, unmap, invalidation, and dereference operations."**

As you can see, the description for CVE-2018-17182 contains terms such as "Linux kernel," "overflows" and "gain privileges," which may suggest a certain level of impact score.

In the absence of any published score(s), Predictive Prioritization predicts a CVSS impact score of 5.5. The predicted threat score at publication is close to 0, which yields an overall Predictive Prioritization score of 5.5.

However, following the publication of a proof-of-concept exploit on Exploit Database on September 26, 2018, the CVE's likelihood of threat starts to increase and the Predictive Prioritization score peaks at 9.5 following discussion on the dark web and social media four days later. The Predictive Prioritization score remains high for the next few weeks until we see a reduction in the CVE's likelihood of threat in response to the changing threat landscape.

On November 28, a full 70 days after the vulnerability was published on NVD, the CVE is assigned a CVSS vector (AV:L/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H), which yields a v3 base score of 7.8. Although the actual v3 impact score (5.9) was slightly higher than the predicted value (5.5), Predictive Prioritization empowers security operations teams to effectively prioritize a vulnerability for remediation during this "dark period" between publication and when the CVSS vector is published.

In short, Predictive Prioritization is more accurate than CVSS alone because:

- Predictive Prioritization rescores every CVE daily to constantly align VPRs with the shifting threat landscape.
- Predictive Prioritization adds sophisticated threat intelligence to predict which vulnerabilities will be exploited in the near-term future, providing a better gauge of actual vs theoretical risk.
- Predictive Prioritization delivers a more granular, differentiated view of vulnerabilities (versus CVSS, which scores the majority as Critical or High).

Comparison of Remediation Strategies

One way to estimate how well a remediation prioritization strategy will perform is to use historical data to calculate its "risk coverage." Think of risk coverage as the percentage of vulnerabilities that should have been remediated, as correctly predicted by the model. For example, if 100 CVEs have threat activity and the model correctly predicts 60 of these CVEs, then the risk coverage is 60%.

The chart below compares the risk coverage for four different remediation strategies where the scores have been filtered for the top 5,000 (highest to lowest):

- VPR
- CVSS*
- CVSS (filtered for CVEs on Exploit Database)
- CVSS (filtered for CVEs on Metasploit**)

Risk Coverage Comparisons

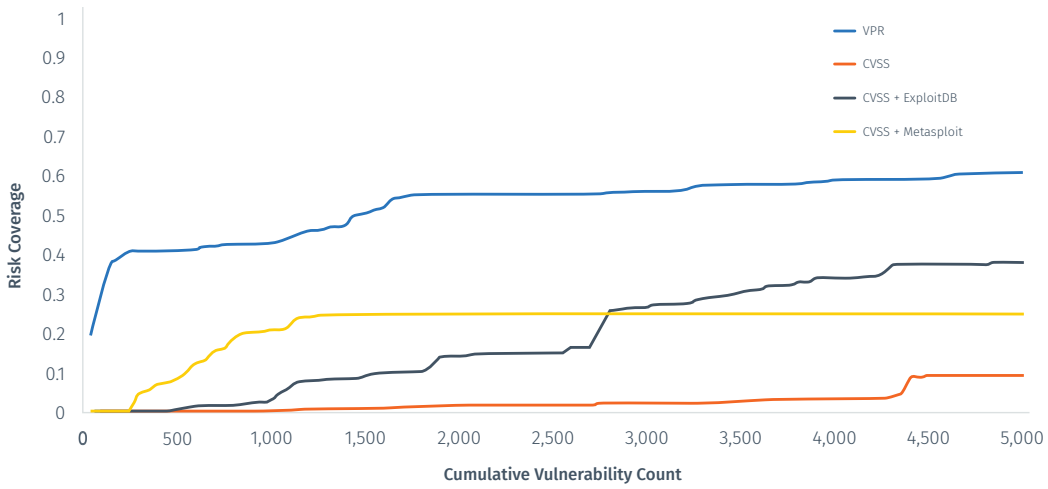

Figure 8. Risk Coverage Comparisons

*CVSS scores used for this chart are a blend of versions 2 and 3. Version 3 is used if available, otherwise version 2 is used.
**Note: There are only 1,284 CVEs on Metasploit in this sample, which is why the line flatlines after this point.

The table below summarizes key data points from the chart:

Risk Coverage

	VPR	CVSS	CVSS + Exploit Database	CVSS + Metasploit
Top 500 CVEs by score	41.3%	0.0%	0.9%	9.0%
Top 1,000 CVEs by score	43.0%	0.4%	3.1%	21.1%
Top 2,000 CVEs by score	55.2%	1.3%	14.3%	25.1%
Top 5,000 CVEs by score	61.0%	9.4%	38.1%	25.1%

VPR stands out as the most efficient scoring strategy, capturing a significantly higher proportion of vulnerabilities with threat activity for the same quantity of vulnerabilities to remediate.

Drive Operational Efficiency with an Automated Approach

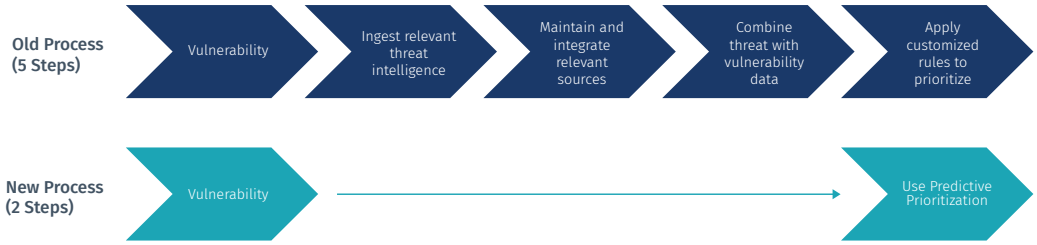

Figure 9. Operational Efficency: Threat and Vulnerability Prioritization

Predictive Prioritization helps organizations achieve vulnerability management operational efficiency in an era when top talent is tough to attract and even more difficult to retain and budgets are tight. That's important because few organizations have mastered today's common five-step process for threat and vulnerability prioritization. With Predictive Prioritization, security teams can focus their attention on the true risks to the business as opposed to theoretical fairytales.

Conclusion

Vulnerability management has become more difficult as the number of vulnerabilities continues to increase. Patching all vulnerabilities isn't practical given limited cybersecurity and IT resources, so organizations must prioritize their vulnerability management efforts.

Most organizations prioritize vulnerabilities using CVSS scores. However, the scoring changes from CVSSv2 to CVSSv3 have resulted in nearly twice the number of high and critical CVEs, which means CVSS scoring alone doesn't solve the problem. There are still too many vulnerabilities to manage.

Predictive Prioritization builds on CVSS scores, adding threat intelligence and machine learning to render VPRs that are more accurate than other remediation strategies. Using Predictive Prioritization, organizations can dramatically improve their remediation efficiency and effectiveness by focusing on the 3% of vulnerabilities that have been or will likely be exploited.

Predictive Prioritization is a key capability within the Cyber Exposure platform, providing security teams with actionable insights to answer the critical question: Where should we prioritize?

It is available today in Tenable.sc (formerly SecurityCenter) on-premises and will be in the cloud-based Tenable.io later in 2019.

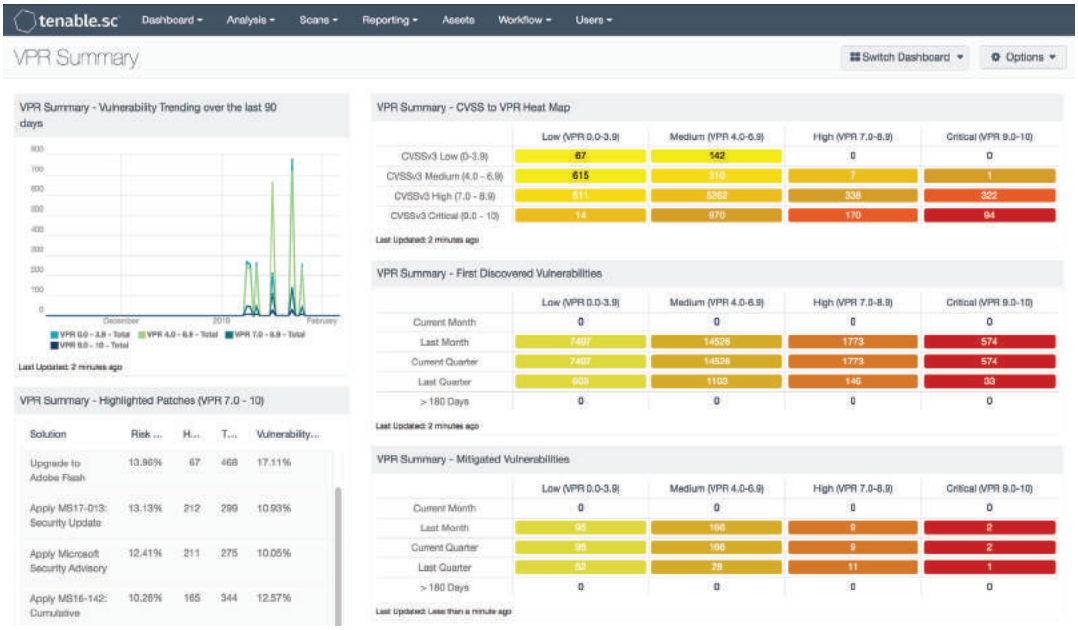

See Predictive Prioritization in action. Request a demo now.

7021 Columbia Gateway Drive
Suite 500
Columbia, MD 21046

North America +1 (410) 872-0555

www.tenable.com

02/10/19 V01