Health-ISAC: Risk-Based Approach to Vulnerability Prioritization

Authors:

Brian Bizon - Organon
David Glosser - Regeneron
Drew Vravick - AbbVie
Erik Dacey - Merck
Tyler Curry - Health-ISAC

Health-ISAC™ Collaborating for Resilience in Healthcare

Contents

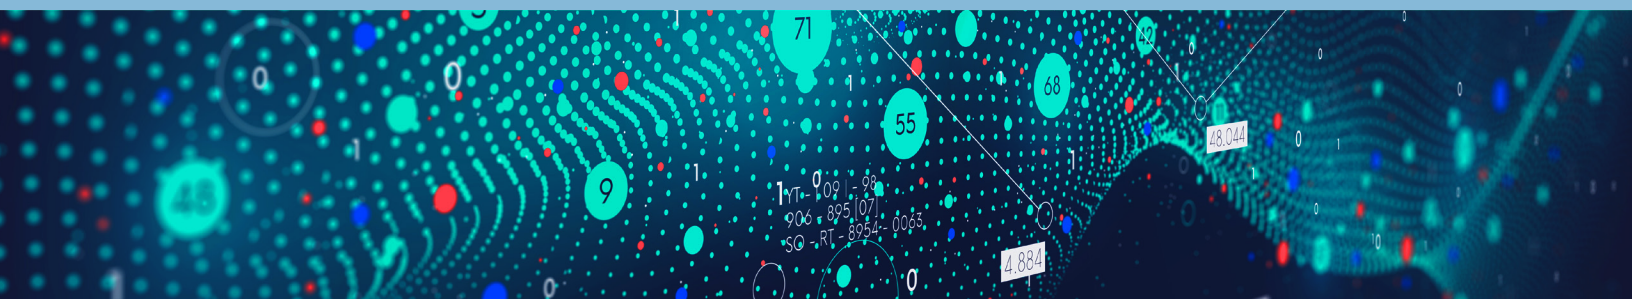

Abstract

With over 15,000 vulnerabilities already identified in 2023 and 25,227 in 2022, organizations are reliant upon the resources available to them.[1] Organizations are increasingly overwhelmed by the volume of findings and the challenging task of triaging vulnerabilities to determine which to address first in a timely and well-judged manner. As a result, there is a need for maturing vulnerability management processes and a shift away from traditional severity ratings. With the evolution of threat actor capabilities greatly influencing the rise of exploitation, it is important for organizations to implement sustainable frameworks and standards for prioritization in vulnerability management. This paper stands as the first iteration of a series of communications regarding vulnerability management, focusing on the importance of prioritization and its applicability to organizations using a variety of recommended concepts.

Executive Summary

Network security teams are often encumbered with the ongoing release of vulnerabilities that are either publicly disclosed or identified as zero-days by vendors and security researchers. Each of these vulnerabilities' severity and exploitability levels is associated with a Common Vulnerability Scoring System (CVSS) score and, often, with a Common Vulnerabilities and Exposures (CVE) number. These swaths of information have proven cumbersome and, at times, can pose a conundrum to organizations concerning their vulnerability management capabilities. Only 2-7 percent of all published vulnerabilities are ever exploited in the wild and, in many cases, are ignored due to a lack of prioritization.[2]

The concept of prioritization in vulnerability management is significant as it helps to support effective mitigation and remediation strategies across different organizational capability levels. The correlation between prioritization and organizations' capability level is closely aligned as it can help security teams communicate effectively with stakeholders, identify asset value, and develop remediation policies conducive to the continuity of business-critical systems. Prioritization is a process that spans all capability levels and allows security teams to properly allocate resources to address vulnerabilities associated with severity levels that exceed the organization's risk appetite.

1 https://www.cvedetails.com/browse-by-date.php
2 https://www.first.org/epss/model

Using Base CVSS Scoring

Depending on the organization's size, staffing constraints, maturing of the security team, or if the vulnerability management program is in its infancy, a company may choose to remediate all critical and high severity vulnerabilities utilizing the scoring system within the chosen vulnerability scanning tool. This is a good baseline for a newly created vulnerability management program.

Exploitability Metrics

Attack Vector (AV)

Network (N) | Adjacent (A) | Local (L) | Physical (P)

Attack Complexity (AC)

Low (L) | High (H)

Privileges Required (PR)

None (N) | Low (L) | High (H)

User Interaction (UI)

None (N) | Required (R)

Scope

Scope (S)

Changed (C) | Unchanged (U)

Impact Metrics

Confidentiality Impact (C)

High (H) | Low (L) | None (N)

Integrity Impact (I)

High (H) | Low (L) | None (N)

Availability Impact (A)

High (H) | Low (L) | None (N)

Base Metric Group

Exploitability metrics	Impact metrics
Attack Vector	Confidentiality Impact
Attack Complexity	Integrity Impact
Privileges Required	Availability Impact
User Interaction	
Scope	

Temporal Metric Group

- Exploit Code Maturity
- Remediation Level
- Report Confidence

Environmental Metric Group

- Modified Base Metric
- Confidentiality Requirement
- Integrity Requirement
- Availability Requirement

This methodology has been employed at companies where security personnel have multiple roles and responsibilities, including vulnerability management, incident response, endpoint security, network security, etc. Focusing on the intricacies of vulnerability management is difficult for small teams. Management may be concerned that an analyst might make a judgment call on prioritizing a specific vulnerability over another and make the wrong choice, leaving the organization susceptible to a specific threat. Using this methodology eliminates the human element from prioritization entirely.

The downside to this methodology is that remediation teams may be overwhelmed by the sheer number of vulnerabilities they are asked to focus on. For example, a knowledgebase search via a commercial vulnerability scanning tool found over 50,000 detections with a severity of four or five (High and Critical). The number of detections for severity levels one, two, and three was over 56,000. ¬Additionally, threat actors may not always exploit the highest severity vulnerabilities and instead chain together multiple exploits of less severe vulnerabilities to gain access to systems.

Focusing on Known Exploited Vulnerabilities

A more risk-based approach would be to focus on known exploited vulnerabilities. On November 3, 2021, CISA released Binding Operational Directive 22-01, which focused on reducing the risk of known exploited vulnerabilities.[3]

Known exploited vulnerabilities should be the top priority for remediation. Based on a study of historical vulnerability data dating back to 2019, less than four percent of all known vulnerabilities have been used by attackers in the wild. Rather than having agencies focus on thousands of vulnerabilities that may never be used in a real-world attack, BOD 22-01 shifts the focus to those vulnerabilities that are active threats. CISA acknowledges that CVSS scoring can still be a part of an organization's vulnerability management efforts, especially with machine-to-machine communication and large-scale automation. It is important to remember that the Directive is intended to help agencies prioritize their remediation work; it does not release them from any compliance obligations, including resolving other vulnerabilities.[4]

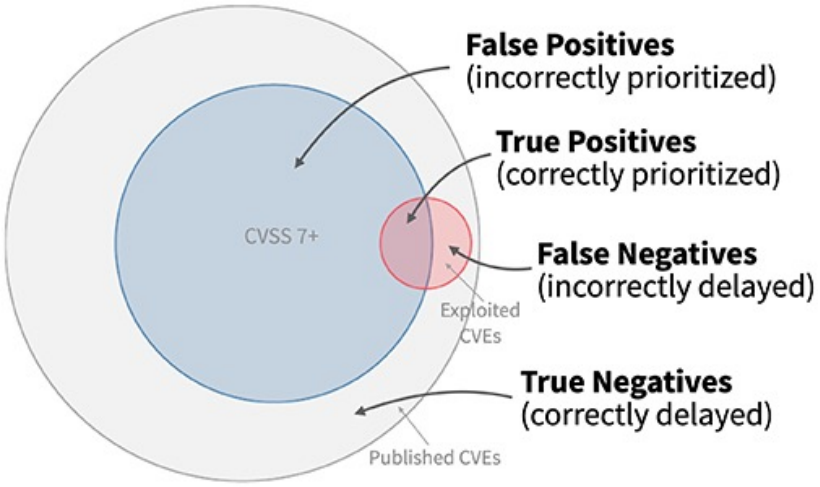

This methodology significantly reduces the number of vulnerabilities that need immediate attention. As of July 13, 2023, there were less than 1,000 vulnerabilities on the list. It also ensures practitioners focus on vulnerabilities that pose the greatest threat to organizations. A process that keeps an organization safe would likely include focusing on CISA's KEV list and pivoting to remediate non-exploited vulnerabilities with critical and high severity levels.

3 https://www.cisa.gov/news-events/directives/bod-22-01-reducing-significant-risk-known-exploited-vulnerabilities
4 https://www.cisa.gov/news-events/directives/bod-22-01-reducing-significant-risk-known-exploited-vulnerabilities

Device Context or Placement

Network location is a critical point of interest that organizations must account for in vulnerability prioritization. With the volume of CVEs discovered and zero days exposed on internet-facing assets, vulnerability management, and technology teams need to understand the functionality and placement of devices in their network. A critical prioritization aspect is understanding and detecting applications and technology running on your perimeter. As vulnerabilities are disclosed, you will want to ensure you have an emerging threat or critical vulnerability response to security flaws that are publicly facing. Internet-facing vulnerabilities and misconfigurations should always be a priority since threat actors have an increased appetite to expose them, and they pose as a path of least resistance, putting organizations at a greater risk of compromise.

For systems that are internally facing or not accessible from the internet, these should fall under an internal service level agreement (SLA) remediation timeline. Organizations should consider making SLAs according to different aspects regarding network location. For instance, internet-facing assets should have a shorter SLA than internally-facing assets. The layout of a network plays a critical role as well when taking into consideration the high volume of vulnerabilities and how vulnerability management teams assess which to focus on first. When looking at internal vulnerabilities, the main focus in prioritization would be lateral movement and how an attacker can pivot internally to systems potentially housing sensitive data. You will want to prioritize vulnerabilities allowing an attacker to gain control of a system or move laterally in the network. Some keys to prioritizing these vulnerabilities will include using vulnerability priority ratings. Most of the tools today are leveraging additional scoring features, like the exploit prediction scoring system (EPSS), to help analysts begin to prioritize vulnerabilities. The vulnerability priority ratings deployed by vendors allow analysts to assess which security flaws must be remediated first based on the likelihood of exploitation within their network. This enables organizations to operate in a manner that is conducive to a risk-based approach.

If there is a remote code execution vulnerability that impacts an organization's web servers, they would likely need to patch them before addressing a similar issue impacting their internal web applications. Likewise, a local privilege escalation vulnerability in an operating system known to be exploited through phishing emails would be a priority to fix on a workstation before a server. By incorporating this location-based context, patching teams will work to remediate vulnerabilities by attack vector, exploitability, and severity.

Asset Value

Asset value is another important factor in vulnerability prioritization. Analysts must know the asset's value as they leverage device context and placement. Teams may begin using a ranking system within their application repository to identify critical assets. The vulnerabilities associated with these critical assets can serve as grounds to be prioritized first when presented to stakeholders and allow analysts to influence decisions to remediate vulnerabilities impacting business-critical assets.

If a device that is of the utmost importance to the operation of the business or holds critical information were to be compromised, it could be catastrophic to the organization. This is why it is recommended to prioritize patching these devices over others. Incorporating business impact into severity weighting gives a more accurate view of risk to the company.

To effectively implement this strategy, an accurate and agreed-upon business impact value per company asset is needed. Ideally, this is centrally located, such as in a Configuration Management Database (CMDB). Although most industry CMDB products provide an asset discovery solution to help maintain inventory accuracy, it will only be partially absolved of challenges.

Compensating Controls

Most organizations have layered security controls or defense-in-depth strategies to mitigate attacks executed by advanced security threats. These security controls should make it more difficult to exploit vulnerabilities. Raising or lowering the severity of vulnerabilities based on compensating controls is controversial. Those responsible for remediation often campaign to lower the severity of vulnerabilities under the assumption that the control is effective. Changing a vulnerability's severity or risk rating without sufficient data to support the change can cause teams to focus on the wrong vulnerabilities and weaken an organization's security posture.

In cases where an organization has personnel with red teaming expertise, it is recommended to test the exploitation of vulnerabilities against the company's security stack in a sandboxed environment. Another option is to use a breach and attack simulation tool to mimic the tactics, techniques, and procedures (TTPs) of the exploitation activities observed in malicious operations. Having this data will help determine if certain vulnerabilities' severity or risk rating can be decreased or increased.

EPSS – Exploit Prediction Scoring System

It is extremely challenging for IT and security teams to follow the patch everything approach – there are just too many vulnerabilities to fix them all in a timely manner. The Common Vulnerability Scanning System (CVSS) is the most frequently cited rating system to assess the severity of security vulnerabilities. However, it has been criticized as not being appropriate to use as a sole method to assess and prioritize risks from vulnerabilities. Also, only a limited subset of published vulnerabilities is ever observed being exploited in the wild.

One technique to prioritize which vulnerabilities should be patched first is to emphasize vulnerabilities known to be exploited and not rely only on the Base CVSS Score, which does not consider if a vulnerability has been exploited. To support this effort, CISA has published a list of Known Exploit Vulnerabilities (commonly called the CISA KEV list). CISA strongly recommends all stakeholders include a requirement to immediately address KEV catalog vulnerabilities as part of their vulnerability management plan. The KEV list can be used to prioritize remediation efforts on the subset of vulnerabilities that are known to be exploited by threat actors.

Another tool that can help prioritize vulnerabilities is called the Exploit Prediction Scoring System (EPSS).[5] EPSS is an open, data-driven effort that uses a machine-learning model to predict the likelihood or probability that a vulnerability will be exploited in the wild to assist defenders in prioritizing vulnerability remediation efforts more effectively. Like CVSS, EPSS is governed by the Forum of Incident Response and Security Teams (FIRST). EPSS uses data from sources like the MITRE CVE list, data about CVEs such as days since publication, and observations from exploitation-in-the-wild activity from security vendors. The EPSS model produces a probability score between zero and one (0 and 100%). The higher the score, the greater the probability that a vulnerability will be exploited.

5 https://www.first.org/epss/model

EPSS Score compared to CVSS Base Score (NVD)

Point density is represented by color, yellow is less dense going through red to a deep purple for the most dense areas labeling a random sample of CVEs with higher values for reference.

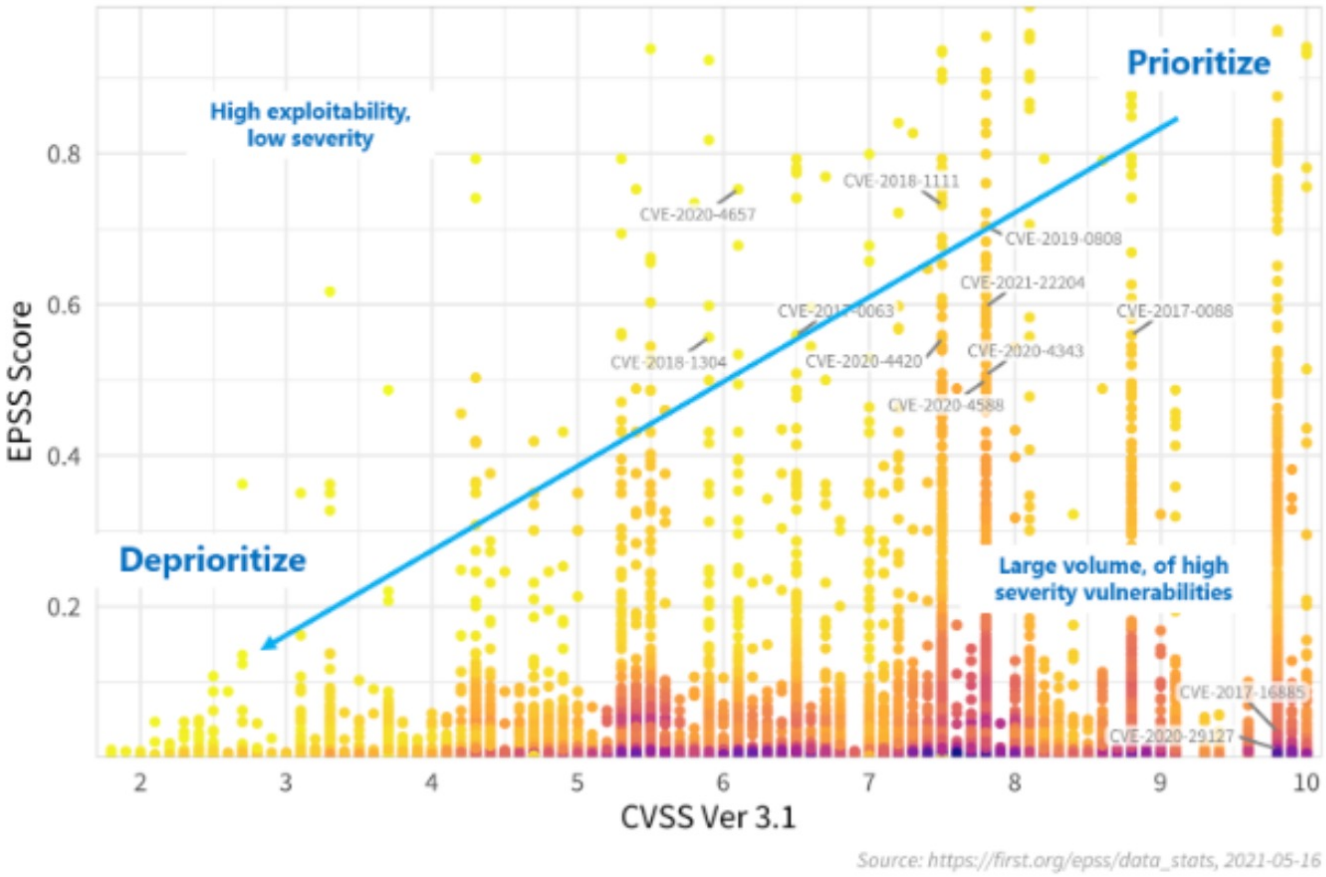

Source: https://first.org/epss/data_stats, 2021-05-16

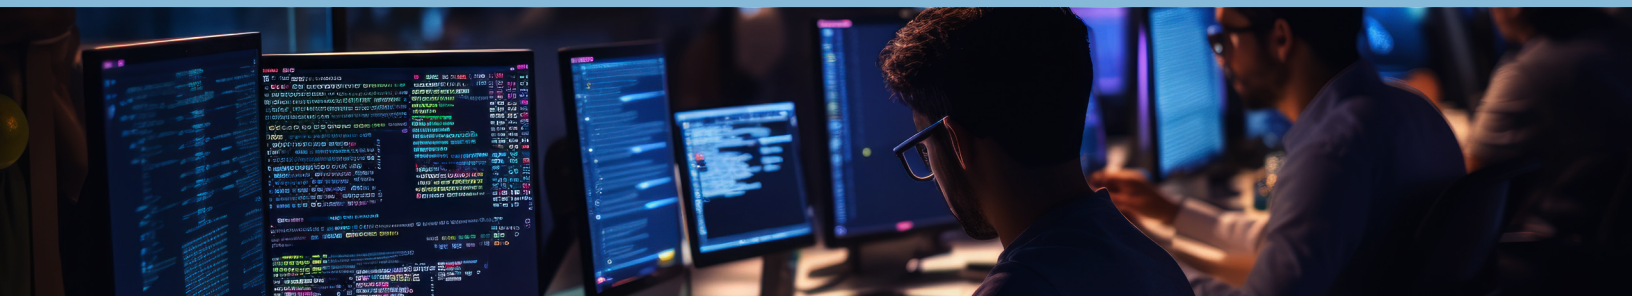

SSVC – Stakeholder-Specific Vulnerability Categorization

While EPSS incorporates current threat information from real-world exploit data to enable defenders to better prioritize vulnerability remediation, stakeholder-specific vulnerability categorization (SSVC) focuses on values, including the security flaw's exploitation status, its impact on safety, and the prevalence of the affected products. SSVC improves vulnerability management processes and accounts for diverse stakeholders. The vulnerability analysis methodology was developed by Carnegie Mellon University's Software Engineering Institute in coordination with the US Cybersecurity and Infrastructure Security Agency (CISA) to operate as a decision tree that allows for flexibility in its application, to which CISA previously shared a guide with slightly modified values.[6] When CISA becomes aware of a vulnerability, the following possible decisions illustrated in Table 1: Vulnerability Decision, Possible Outcomes, are made[7]:

- **Track:** Currently, the vulnerability does not require action in which the organization would continue to track the vulnerability and reassess it if new information becomes available; CISA recommends remediating Track vulnerabilities within standard update timelines.
- **Track*:** The vulnerability contains specific characteristics that may require closer monitoring for changes; CISA recommends remediating Track* vulnerabilities within standard update timelines.
- **Attend:** The vulnerability requires attention from the organization's internal, supervisory-level individuals, in which necessary actions include requesting assistance or information about the vulnerability and may involve publishing a notification either internally and/or externally; CISA recommends remediating Attend vulnerabilities sooner than standard update timelines.
- **Act:** The vulnerability requires attention from the organization's internal, supervisory-level, and leadership-level individuals, in which necessary actions include requesting assistance or information about the vulnerability, as well as publishing a notification either internally and/or externally. Typically, internal groups would meet to determine the overall response and then execute agreed upon actions; CISA recommends remediating Act vulnerabilities as soon as possible.

Implementing SSVC allows organizations to leverage its customizability to help analysts decide on vulnerability response actions consistent with maintaining the confidentiality, integrity, and availability of enterprise systems as agreed upon with leadership. SSVC aims to be a dynamically applied concept as new versions are released to recognize improvements with the addition of the coordinator stakeholder perspective, updates to terminology, integration of feedback on decision point definitions, and tools to support practical use.[8] With that, it is safe to say that SSVC focuses on more than just base scores as a stand-alone prioritization method.

6 https://www.cisa.gov/sites/default/files/publications/cisa-ssvc-guide%20508c.pdf
7 https://www.cisa.gov/stakeholder-specific-vulnerability-categorization-ssvc#Cisa's%20SSVC%20Calculator
8 https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=653459

Leveraging SSVC helps analysts make decisions that address a particular issue or subset of issues over another. Organizations can efficiently prioritize and triage vulnerabilities while traversing the uncertainties of what issues to address first. Every decision is the result of a logical combination of triggers set by leadership in response to factors that may include the vulnerability's state of exploitation (i.e., proof of concept, active exploitation), the level of difficulty for an adversary to exploit it, and its impact to public safety. The factors in this example are non-exhaustive; however, an analyst would then collect evidence of the relevant triggers and utilize the logic of the decision tree to establish triage priority decisions.

Once the decision tree's logic is used to establish triage priority decisions, organizations can apply one of the five values – track, track*, attend, and act. Outcomes including evidence of exploitation, malware propagation, technical impact, mission prevalence, and public well-being influence which value is chosen to address the vulnerability. CISA uses the following decision points and associated values for making vulnerability scoring decisions:

- **State of Exploitation Evidence:** This outcome is straightforward and is often determined shortly after or around the time the vulnerability is disclosed. Oftentimes, this is observed, especially in cases where a zero-day is discovered in widely used software or systems. When gathering information regarding this decision point, analysts must determine whether a proof of concept is publicly available and/or if the vulnerability is being actively exploited by threat actors.
- **Automatable:** When considering this outcome, analysts must determine if the vulnerability is self-propagating. In situations where it is not, this would indicate that there are barriers in place that a threat actor would need to circumvent. These layers include security controls that may include address space layout randomization (ASLR) user interaction, enforcement of required privileges, or local/physical access to the network.
- **Technical Impact:** This outcome falls under two categories, including partial or total. In the case of total technical impact, a threat actor would compromise a vulnerability in which an analyst would need to determine if they would have complete control over the affected product (i.e., remote code execution, privilege escalation, stolen information, or exposed credentials). For partial technical impact, this would include whether there was just a denial of service or partial information disclosure.
- **Mission Prevalence:** This outcome assesses the number of affected products in a particular environment while focusing on the product's functionality. The category effectively allows analysts to determine if the product has a dominant presence in its environment and whether it impacts the continuity of operations. This factor also considers the level of support necessary to mitigate related issues.
- **Public Well-Being:** This outcome focuses on the financial, social, and psychological impacts of exploitation that a vulnerability would have in an organization's environment. Analysts would need to consider several different issues, such as whether impacts on product users incur irreversible, material, or minimal damages. These damages may include physical harm, exposed payment information, or negative reputational implications.

A practical use case in which this methodology can be applied is a vulnerability management team's prioritization response to the Citrix ShareFile vulnerability, identified as CVE-2023-24489. Leveraging triggers agreed upon by leadership, an organization would likely choose the act value after running information collected by analysts against the decision points and associated values previously mentioned. In this scenario, the decision to choose the act value is influenced by the existence of proof-of-concept (PoC) code in tandem with evidence indicative of targeted attacks and in-the-wild exploitation. The vulnerability is not self-propagating; however, the technical impact falls under the total category, as unauthenticated attackers can upload files and arbitrarily execute code on all compromised user-managed installations.

Like other enterprise-grade file-sharing applications, including Accellion File Transfer Appliance (FTA) devices, the GoAnywhere MFT platform, and the MOVEit Transfer solution, the mission prevalence of Citrix ShareFile is prominent as it likely has a dominant presence in its environment and can impact business operations. Additionally, the potential exposure of information contained within vulnerable Citrix ShareFile instances could prove to be catastrophic to public well-being. Previous exposures to enterprise-grade file-sharing applications contained sensitive data, including but not limited to medical, personally identifiable information (PII), and financial details.

Policy and Governance

Regardless of the size of the organization and the resources available, it is essential to have the vulnerability management program documented in a company policy. The policy document should make clear the organization's expectations for addressing vulnerabilities. For example, a program must be in place to manage vulnerabilities, define the role that is to be held accountable for running the program, specify who has ownership, and who is responsible for remediation. The policy should be reviewed and approved by the organization's management, indicating their support of the program, and should be made available across the company. Many organizations have policy awareness programs that require annual review of their policies through mandatory training. Many policies will include guidance on vulnerability management remediation timeframes based on the risk to the company. Timeframes can be set based on location in the enterprise, business impact, and known exploitable vulnerabilities, as previously discussed.

The document should also include clear guidance on what is within the scope of the policy. An example could be that the policy is applied globally to product or application owners accountable for identifying and remediating vulnerabilities on company-managed devices connected to the enterprise network. There should also be provisions for devices that are not company-managed (i.e., Software-as-a-Service or SaaS) and what requirements these devices must meet. These requirements can be included in the contract language during the initial engagement with the SaaS vendor. Finally, a deviation process should be implemented to review and approve cases where the policy statement could not be followed. The deviation should include the person accountable for remediating the change, their plan of action, expected timeframes, and any available compensating controls.

Metrics

Metrics are a vital part of any vulnerability management program. They not only highlight the effectiveness of the program but identify areas that need improvement. Focusing on the number of critical, high, medium, and low severity vulnerabilities in an environment is insufficient to determine that remediation efforts are meeting goals. Compartmentalizing metrics by technology, placement on the network, and the SLA outlined in the company policy are ways to identify areas of improvement. Distinguishing known exploited vulnerabilities from those that are not currently exploited can reduce noise and point teams to remediation efforts that need more visibility.

Organizations should also focus their metrics on key risk indicators versus key performance indicators and highlight specific insights obtained from their vulnerability data. For instance, an organization may associate the Chrome platform as having a high level of risk because the time to remediate the browser vulnerabilities is excessive. This may be evident in that, on average, most of the vulnerabilities on the Chrome platform remain unresolved after 90 days, as opposed to Edge mostly being remediated to completion after 30 days. As a result, organizations would need to shift their attention to Chrome. Through this practical concept, organizations can use performance metrics to show areas of risk versus tracking individual vulnerabilities and allow space for actionable insights.

Conclusion

Vulnerability management professionals can no longer fully rely on the data that comes directly from a scanning tool to determine the best way to secure organizations. To be effective, teams must clearly understand their environment and combine that context with data from multiple sources to reduce the risk posed by vulnerabilities. There is no one-size-fits-all method to perform vulnerability prioritization. Teams must determine what works best for their organization based on several different criteria. It is recommended to combine multiple prioritization methodologies identified in this whitepaper to find what fits best in your organization.

References:

1. **CVE Details:** Browse Vulnerabilities by Date

2. **FIRST:** The EPSS Model

3. **BOD 22-01:** Reducing the Significant Risk of Known Exploited Vulnerabilities

4. **BOD 22-01:** Reducing the Significant Risk of Known Exploited Vulnerabilities

5. **FIRST:** The EPSS Model

6. **CISA:** Stakeholder-Specific Vulnerability Categorization Guide

7. **CISA:** Stakeholder-Specific Vulnerability Categorization

8. **Prioritizing Vulnerability Response:** A Stakeholder-Specific Vulnerability Categorization (Version 2.0)