Targeted Data Visualization and Reporting Approaches for Vulnerability Management at Enterprise Organizations

Ryan Painter
painter@rmu.edu
Computer and Information Systems
Robert Morris University
Moon Township, PA 15108 USA

Abstract

Large enterprise organizations operating a vulnerability management program are inundated with large quantities of data resulting from vulnerability scans conducted against organizational technology assets. In order to ensure this information can be appropriately interpreted by the various teams responsible for operating and managing the vulnerability management program, targeted data visualization and reporting approaches are key. This paper outlines a targeted data visualization and reporting approach for three key organizational audiences: vulnerability remediation teams, senior or middle management, and executive leadership.

Keywords: Vulnerability management, vulnerability remediation governance, data visualization, reporting

1. INTRODUCTION

Organizations of all sizes continually find themselves targeted by malicious cyber threat actors, resulting in the potential for serious negative financial consequences in the event of a successful attack. In the first half of 2021, a harrowing record was reached for the average ransomware payment amount made to threat actors by organizations, reaching $570,000 from 2020's average of $312,000 (Baylor et al., 2021). To effectively manage the risk of negative organizational impact from cyber threat actors, large organizations employ numerous strategies to defend organizational technology resources, monitor for successful attacks, and reduce the attack surface of their assets.

Vulnerability management serves as one pillar of a cyber defense strategy, wherein organizations systematically identify vulnerabilities present on their networks at scale through scanning or through penetration testing engagements, report on identified findings, and either track through to remediation or manage risk through exception management and/or the application of compensating controls. As penetration testing relies on human resources to conduct vulnerability assessments in a typically manual manner, vulnerability scanning is more often used to depict a fuller picture of an organization's vulnerability posture across all assets which can be scanned. As one could imagine, organizations with employees in the hundreds of thousands face not only the significant challenge of architecting their vulnerability scanning solution in a manner that is able to effectively scan each of their vast numbers of systems but also the feat of analyzing data generated from these scans to support vulnerability remediation activities and best mitigate organizational cybersecurity risk.

To assist with the challenge of interpreting, communicating, and prioritizing vulnerability data, this paper serves to provide a vulnerability management data visualization and reporting strategy for cybersecurity practitioners at large enterprise organizations. First, a background of

vulnerabilities and corresponding metadata are discussed to provide a foundation for the nature of data that is collected through the vulnerability scanning process. Next, data visualization concepts and best practices are discussed to provide perspective on how the paper's proposed approach was developed. Then, the proposed data visualization and reporting approach is outlined, along with concluding remarks.

2. VULNERABILITY DATA AND SCANNING

The National Vulnerability Database (NVD), operated by the US government via the National Institute of Standards and Technology, documents individual vulnerabilities as they are discovered by security researchers and technology vendors over time. As part of these records, numerous attributes are recorded pertaining to vulnerabilities, including a Common Vulnerabilities and Exposures (CVE) ID; Description; Severity; References to Security Advisories, Solutions, and Tools; Weakness Enumeration; and Known Affected Software Configurations (National Institute of Standards and Technology [NIST], n.d.).

The CVE ID serves as a unique identifier for a given vulnerability. The Description field provides a summary of the nature of the vulnerability, as well as some high-level information on the affected product(s). Severity incorporates the Common Vulnerability Scoring System (CVSS) score, versions 3.x and 2.0, to provide a numerical representation of the inherent risk associated with the vulnerability (MITRE, 2021). Discussion of the specific means through which a CVSS score is calculated is outside the scope of this paper, but an overview of the data points which feed into this calculation will follow.

The References to Security Advisories, Solutions, and Tools field includes information from vendors which develop the affected product, or from security advisories published by external security firms, for example, which pertain to the vulnerability. The Weakness Enumeration pertains to any Common Weakness Enumeration (CWE) IDs associated to the vulnerability, which would indicate any CWE common software security weaknesses present in the affected product which contributes to the vulnerability in question (MITRE, 2021). Last, Known Affected Software Configurations details the specific Common Platform Enumeration (CPE) values that are affected by the vulnerability; CPE provides a means to describe a specific technology product type and corresponding version number in a standardized manner (NIST, 2021).

Common Vulnerability Scoring System (CVSS)

Multiple vulnerability characteristics are used to calculate a CVSS score, each of which factor in some manner into the inherent risk posed by the presence of a given vulnerability. Although multiple versions of the CVSS specification are actively used by organizations today, this paper will limit its focus to CVSS version 3.1 for simplicity. Within CVSS 3.1, attributes are included in one of three different categories: base metrics, temporal metrics, and environmental metrics. Within the base metric group are exploitability metrics (attack vector, attack complexity, privileges required, and user interaction), impact metrics (confidentiality impact, integrity impact, and availability impact) and the scope metric. Base metrics are generally considered to be more static in nature, relating to the explicit characteristics of a vulnerability that would not change over time (FIRST, 2019).

Temporal metrics, on the other hand, which include exploit code maturity, remediation level, and report confidence, are expected to shift over time and impact the value of the CVSS score accordingly. As an example, when a vulnerability is announced, exploit code might not be publicly available, meaning that vulnerability would probably be less likely to be exploited by threat actors who do not have the requisite knowledge to develop their own exploit code. Were working exploit code to be made public, however, the population of individuals who have the capability to exploit the vulnerability would likely increase, increasing the risk associated with the vulnerability being present on a system or within an application. Temporal metrics are intended to capture some of these changing factors over time as conditions change.

Last, the environmental metrics group intends to reflect the control environment of a specific individual/organization's technology environment, as the protections put in place to mitigate the risk of certain vulnerabilities and attacker tactics and techniques are likely to vary from environment to environment (FIRST, 2019). Figure 1 details each of metrics groups in CVSS 3.1.

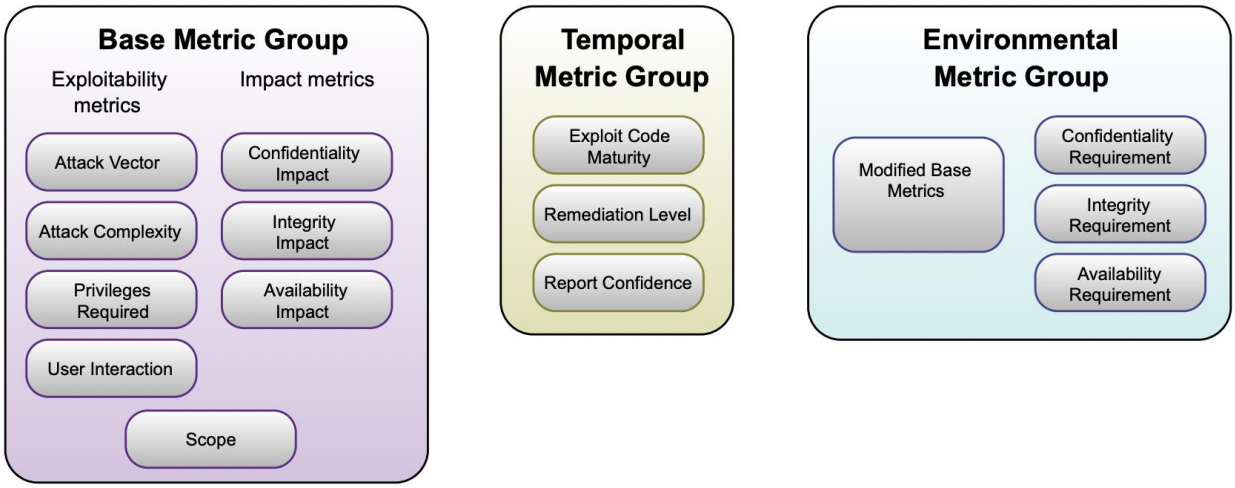

Figure 1. Metric groups in CVSS 3.1 (FIRST, 2019, p. 3)

Vulnerability Scan Data

Vulnerability scanners operate, in part, through the use of plugins that contain logic to detect the presence of specific vulnerabilities (Holm et al., 2011). These plugins are commonly included within a searchable database available to a user and include contextual information pertaining to the vulnerabilities detected by a given plugin. As an example, the Qualys Vulnerability Management product provides several different attributes within its vulnerability plugin database, referred to by Qualys as the KnowledgeBase. Some examples of these attributes include the QID (a unique identifier for each vulnerability plugin), Severity (a value provided by Qualys between one and five indicating the risk associated with the vulnerability detected by the plugin), CVE ID, CVSS Score, Solution (instructions for remediating the vulnerability, commonly through the application of a patch or implementation of a configuration change), Category (broad categorization by Qualys of the product family associated to the vulnerability, such as Windows or Linux), and so on (Qualys, n.d.).

In addition to the fields included in vulnerability scan plugin databases such as the Qualys KnowledgeBase, additional data is generated following the completion of a vulnerability scan against a technology asset. Following such a scan, results indicating whether a targeted asset was able to be scanned successfully, as well as data on the vulnerabilities identified on the target system, the location on the target system where the vulnerable component exists (e.g., an outdated version of a piece of software being located at a specific file path) and other potentially relevant information such as open services and ports, are made available in a scan report. Scan reports can typically be exported from the vulnerability scanning platform in XML or comma separated formats, viewed directly within the vulnerability scanning platform or accessed programmatically via the product's application programmer interface (API). Due to the large volume of vulnerability data generated by large enterprise organizations, this data is often accessed through the API and brought into another platform for analysis. Aside from identified vulnerabilities themselves, information from scan reports which can commonly be of use to security teams include target host IP address, target hostname, target port status, target operating system CPE, installed software, and vulnerability detection dates (Qualys, n.d.).

Other Relevant Data

Thus far, the data presented has focused on the identification phase of the vulnerability management process on the characteristics of the vulnerabilities themselves. However, to aid the remediation phase of vulnerability management after vulnerabilities have been successfully identified, additional data points are often required. In large enterprise organizations, multiple technology teams, each with responsibility for a specific subset of the organization's systems or applications, are

typically in place. For example, distinct teams may be responsible for managing the organization's end user workstation footprint and the organization's server footprint, or perhaps a different team is responsible for the management of on-premise servers versus the management of public cloud-hosted servers. With this split in responsibility, it is common that responsibility for vulnerability remediation exists amongst several distinct groups (although a smaller number of organizations have chosen to operate a specialized, separate vulnerability remediation team with central responsibility for all remediation activities).

Configuration management databases (CMDBs) are central systems where asset data is maintained on technology assets in use within an organization. For a given system, attributes such as (but not limited to) the following may be collected and regularly kept up to date as changes are made: system name, Internet Protocol (IP) address, Domain Name System (DNS) entry, operating system type, operating system version, installed software, physical location, data center rack location, line of business alignment (i.e., certain servers or applications may support a specific line of business within the organization). These attributes, whether business-focused or technology-focused, can identify the appropriate internal team which has responsibility for remediating vulnerabilities on a given system.

3. VULNERABILITY DATA VISUALIZATION AND REPORTING APPROACH

Data visualization, also known as information visualization, refers to an approach to visually depict and facilitate an understanding of data and information (Sharda et al., 2018). As large enterprise organizations find themselves presented with massive quantities of data, whether relating to business operations or to the management of their technology infrastructure, data visualization can serve as a means to enable decision makers to make better sense of this data.

Although vulnerability scanning platforms typically provide out of the box reporting and data visualization capabilities, these capabilities can, at times, not function optimally when presented with the massive volume of vulnerability data that a large enterprise organization is likely to maintain if all or nearly all technology assets in use are in scope for vulnerability scanning. As such, although out-of-the-box reporting features may be sufficient for ad-hoc, limited scope reporting, larger organizations may choose to bring vulnerability data into a separate data analysis platform for recurring analysis and reporting.

As different teams within an organization likely have differing responsibilities pertaining to the vulnerability management program, having varying reporting and visualization options available can help to meet the needs of these groups. For this paper, three audience categories, along with corresponding reporting and visualization approaches, are provided: remediation teams, executive teams, and senior/middle management.

Remediation Team Reporting Remediation teams are those responsible for performing activities to remediate vulnerabilities identified on organization systems and/or applications. Following the identification of vulnerabilities, enterprise organizations commonly impose service-level-agreements (SLAs) pertaining to the deadlines by which a vulnerability of a given severity level is expected to be remediated. For example, a critical severity vulnerability identified on an Internet-facing system, due to the inherent risk of the vulnerability coupled with the fact that the threat of exploitation is higher given that the system faces the Internet, may have an SLA of 24 hours. A low severity vulnerability on an internal system, segmented from the rest of the internal network, with a slew of other compensating controls in place, may have an SLA of 90 or more days, or perhaps may not have an organizational requirement for remediation at all. Regardless, remediation teams are responsible for monitoring for vulnerabilities that have been identified on systems under their management, understanding remediation SLAs, and taking action to ensure that vulnerabilities are remediated by the dates specified by organizational SLAs.

As remediation teams have a primary focus on operations, reporting for these teams should provide forecasting of remediation work due within specific time intervals; as an example, reports detailing vulnerability findings which are nearing remediation SLAs within the next seven, fourteen, or thirty days (or other date intervals) can assist with helping remediation teams and their managers to appropriately prioritize work over time. These reports should contain the data attributes discussed in Section 2. Vulnerability Data and Scanning that pertain to characteristics of the vulnerability itself, as well as where the vulnerability was identified on a target system to

assist remediation teams with updating affected components where necessary. Severity information, whether aligned to the CVSS score of the vulnerability, the severity value provided by the vulnerability scanning vendor, or a custom score calculated by the organization should also be included in remediation team reports to assist these teams with performing remediation in a risk-prioritized manner.

From a data visualization perspective, remediation teams can find use in analyzing the ecosystem of vulnerabilities identified on their systems at a given point in time to identify whether vulnerabilities can be remediated in a more comprehensive way, versus individually in a piecemeal approach as they are identified. For example, through use of a pie chart data visualizations that represents the population of vulnerabilities organized by affected product, a remediation team may identify that seventy five percent of active vulnerabilities on their systems pertain to the Google Chrome web browser. With a trending data visualization that shows the makeup of vulnerabilities identified that is organized by affected product over time, perhaps the remediation team identifies that Google Chrome has continually been the product with vulnerabilities identified in the highest quantity over time. If the organization does not proactively apply updates to Google Chrome as they are released and instead waits for a vulnerability report to be issued following a scan, the team may choose to implement a more proactive, regular patching schedule for this product to reduce the effort expended on responding to vulnerability tickets in the future. Having such data visualizations available to remediation teams can result in these sorts of trends more apparent and allow teams to implement structural improvements to upstream patch and configuration management processes and, in the end, more efficiently focus vulnerability remediation efforts on other items.

Senior/Middle Management Reporting

Senior or middle management employees typically hold responsibility for oversight of the remediation teams responsible for conducting vulnerability remediation and/or of teams aligned to support specific technology products (e.g., Windows servers, Linux servers, Oracle databases, etc.) or lines of business. As such, although these individuals are not responsible for conducting remediation activities themselves, there still exists a need to understand the performance of teams under their management as it pertains to engagement with the vulnerability management process.

Analysis of trend information was discussed in the previous section. As senior or middle management holds responsibility for implementing changes to the operational processes of one or more departments, management would benefit as well from consuming data visualizations that allow them to identify trends relating to the number of vulnerabilities affecting certain products, the remediation solution types for identified vulnerabilities, the total vulnerability risk level of system aligned to a particular line of business, etc. To illustrate this point, consider the following examples. An organization maintains a data visualization and corresponding report which details the remediation solution types for identified vulnerabilities across its technology footprint. Within this data visualization and report, senior and middle management has the ability to filter by line of business, platform type, and so on. By having this data visualization available, management can investigate whether gaps in its patch and configuration management toolset across specific lines of business or subsidiaries are contributing to higher vulnerability counts and move to remedy these gaps where needed.

As another example, consider data visualizations and reports which detail outstanding vulnerability findings which have breached remediation SLA, resulting in non-compliance with organizational policies pertaining to vulnerability management. Management-level awareness of such non-compliance could assist with driving action to remedy these outstanding findings, as well as to perform root cause analysis surrounding what led to this non-compliance to occur in the first place. If a recurring systemic issue pertaining to technology or process is identified over time that continually contributes to remediation non-compliance, management can move to address this issue. Overall, management benefits from access to vulnerability reporting which contributes to their ability to support operational management and strategic decisions for the betterment of the vulnerability management program.

Executive Team Reporting

Executive leadership, although not tasked with managing the day-to-day operations of an organization’s cybersecurity unit, is ultimately accountable for the success of the organization’s cybersecurity program. As such, enterprise leadership is responsible for the establishment of organizational security governance objectives by which cybersecurity management then manages the organization’s cybersecurity program

(Information Systems Audit and Control Association [ISACA], 2015). Because of this, in-depth operational metrics are not likely to be as relevant to executive audiences; however, where vulnerability-oriented reporting can be delivered in a manner that reports performance against the organization's security governance objectives, such reporting could be of use to an executive audience. In some cases, though, especially in situations where a specific vulnerability has gained extensive media attention and client or partner organizations have inquired as to an organization's remediation progress for that vulnerability, executive leadership could express interest in operational reporting pertaining to that remediation effort. Ultimately, like senior or middle management but at a higher level, decision support-oriented data visualizations and reporting will be helpful for an executive audience.

4. CONCLUSION

Large enterprises are faced with high quantities of vulnerability data with a need to prioritize limited resources to direct remediation efforts in a risk-based fashion. By implementing data visualization and reporting approaches aligned to the needs of specific stakeholders in the vulnerability management program, said stakeholders will be better positioned to contribute to the vulnerability management process.

5. REFERENCES

Baylor, R., Brown, J., & Martineau, J. (2021, August 9). *Extortion payments hit new records as ransomware crisis intensifies*. Palo Alto Networks Blog. https://www.paloaltonetworks.com/blog/2021/08/ransomware-crisis/

FIRST. (2019) *Common Vulnerability Scoring System v3.1: specification document* (CVSS Standard v3.1 Revision 1). https://www.first.org/cvss/v3-1/cvss-v31-specification_r1.pdf

Holm, H., Sommestad, T., Almroth, J., & Persson, M. (2011). A quantitative evaluation of vulnerability scanning. *Information Management & Computer Security*, 19(4), 231–247. https://doi.org/10.1108/09685221111173058

Information Systems Audit and Control Association. (2015). *CRISC review manual* (6th ed.). ISACA.

MITRE. (2021, March 13). *CWE - About CWE*. Mitre.org. https://cwe.mitre.org/about/index.html

National Institute of Standards and Technology. (n.d.). *NVD - vulnerability detail pages*. Nvd.nist.gov. Retrieved August 26, 2021, from https://nvd.nist.gov/vuln/vulnerability-detail-pages

National Institute of Standards and Technology. (2021, July 16). *Common Platform Enumeration (CPE) - Security Content Automation Protocol*. Computer Security Resource Center. https://csrc.nist.gov/projects/security-content-automation-protocol/specifications/cpe

Qualys. (n.d.). *Qualys - KnowledgeBase. QualysGuard. Retrieved August 30, 2021, from https://qualysguard.qg3.apps.qualys.com/fo/tools/kbase.php*

Sharda, R., Dursun Delen, & Efraim Turban. (2018). *Business intelligence, analytics, and data science: A managerial perspective*. Pearson Education, Inc