WHITEPAPER

How to avoid a paper tiger vulnerability management program

VULCAN.

Table of contents

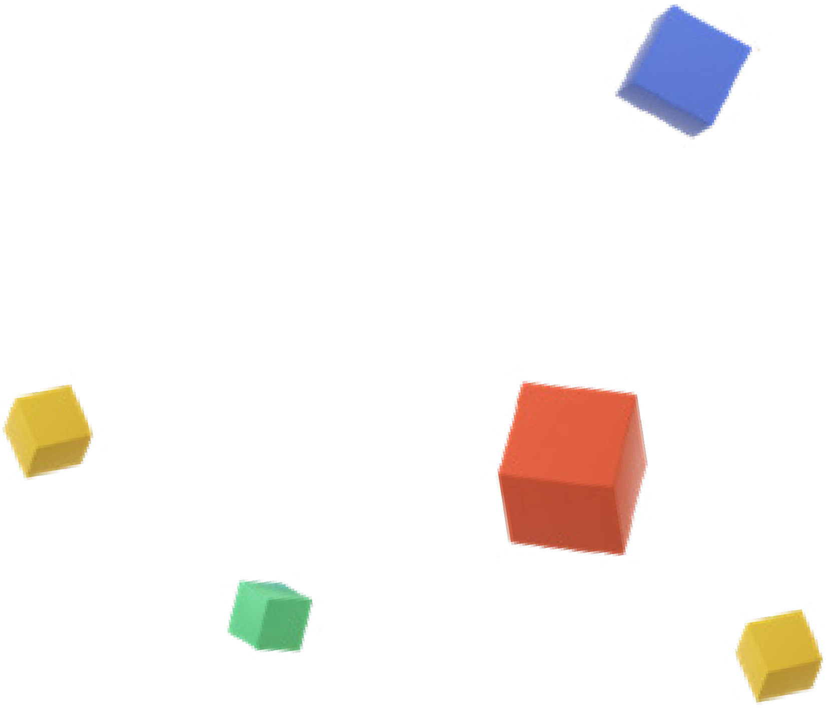

Introduction

There are three main types of vulnerabilities against which organizations must be vigilant: active exploits that represent an immediate and significant risk to the organization's security posture, vulnerabilities that can undermine the organization's compliance posture over time, and vulnerabilities in business-critical products that are widely used throughout an organization. Although each vulnerability type requires its own remediation campaign strategy, these strategies must always align with the organization's security and governance policies as well as its technology and business barriers.

Each remediation campaign strategy has its own challenges. A riskdriven remediation campaign strategy in response to active exploits, for example, is highly time-sensitive. A compliance-driven remediation campaign strategy must achieve high levels of continuous assessment. Product-driven remediation campaigns face many challenges related to technology and business complexity.

In this white paper, we explore the three campaign strategies that can help organizations improve the efficiency of their vulnerability remediation programs. For each strategy, we discuss its use cases and challenges as well as its suitability for different sizes of organizations. We also provide guidelines for running the campaign as effectively as possible.

Urgent vulnerability-driven campaign

The CVSS (Common Vulnerability Scoring System) score range of 7.0 to 10.0 is reserved for vulnerabilities with high technical severity that may have critical business impact. These are usually exploited by threat actors and trending in the news. Some recent examples of high-profile vulnerabilities include:

- BlueKeep: First reported in May 2019, this vulnerability in Microsoft's Remote Desktop Protocol (RDP) makes it possible to execute code remotely using self-propagating worms. In response to indications of proof of concept codes for exploiting the vulnerability, Microsoft issued security patches.

- WannaCry and NotPetya: These two attacks in May and June 2017 exploited the EternalBlue hack of the Windows operating system to encrypt files (WannaCry) or prevent booting (NotPetya) until the victim paid a Bitcoin ransom. Although Microsoft had issued a WannaCry patch two months earlier, the exploit affected 230,000 computers in 150 countries and resulted in about $4 billion in losses. NotPetya affected corporations and government agencies around the globe, causing more than $10 billion in damages.

- Meltdown and Spectre: In early 2018, it was revealed that performance-enhancing features in virtually all CPU chips manufactured over the last 20 years could be exploited to access protected data. The Meltdown variant "melts" hardware-enforced security boundaries, while the two Spectre variants can force programs to reveal secret data.

All organizations, no matter their size, could find themselves running a vulnerability-driven campaign that seeks to remediate a high-impact vulnerability as quickly as possible across an entire environment. The main challenge of this campaign strategy is achieving minimal time to remediation, often with the added pressure of close scrutiny by top management. In large organizations, a company-wide "war room" should be set up to coordinate communications and activities throughout the campaign.

Despite this sense of urgency, the first step in a vulnerability-driven remediation campaign strategy is to gather information based on questions such as:

- Is the vulnerability being actively exploited in the wild? If not, you have a little more time to plan and implement the campaign.

- Which specific systems and versions does the vulnerability target?
Perhaps you will discover that you are less exposed than you thought.

- When will a patch be available and are there known, potential deployment risks? The first Meltdown and Spectre patches, for example, were known to cause performance problems on certain servers. Organizations that could safely do so waited for later patch versions.

- Have the software vendors published workarounds or compensating controls? Just recently Microsoft issued a warning for a serious vulnerability in the Adobe Type Manager library for which no patch is currently available. As a partial workaround, Microsoft advises to disable panes in Windows Explorer to prevent malicious files from being viewed.

Armed with this information, you can now more accurately assess the vulnerability's impact on your specific environment. Where applicable, enable your detection tools to help you map which assets are vulnerable and categorize those assets by their level of criticality. Wherever possible, apply workarounds or compensating controls—either instead of patching or as a first response in order to buy time for a less urgent patch deployment that mitigates operational risk.

For example, for the BlueKeep vulnerability, the NSA recommended that organizations not using RDP should simply disable Remote Desktop Services and its associated port. For those using RDP, the recommended compensating controls were to require Network Level Authentication or two-factor authentication for RDP, or make RDP accessible only via a VPN.

Where patching is required, start with the asset groups that have the highest cyber security risk but the lowest operational risk.
In a fascinating article about how the Goldman Sachs vulnerability management team responded to Spectre and Meltdown, they describe their triage process and how it led them to the decision that employee endpoints (desktops, laptops, mobile devices) should be patched first. Their reasoning was that untrusted code propagated via email attachments or malicious websites would more likely be run on these assets than on servers. In addition, it was known that the performance impact of the patch was more manageable on desktops and laptops than on servers. Given their higher cyber security risk and lower operational risk, it made sense to begin patch deployment on these assets.

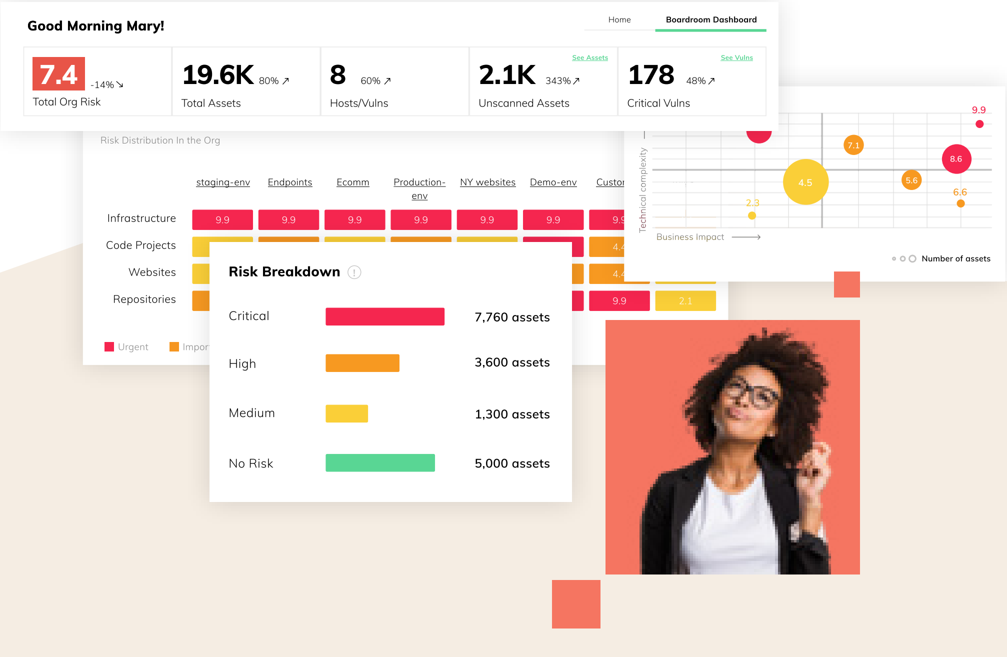
Good Morning Mary!

7.4 -14% ↘	19.6K 80% ↗	8 60% ↗	2.1K 343% ↗	178 48% ↗
Total Org Risk	Total Assets	Hosts/Vulns	Unscanned Assets	Critical Vulns

Risk Breakdown

Critical	7,760 assets
High	3,600 assets
Medium	1,300 assets
No Risk	5,000 assets

Ongoing policy-driven campaign

Many organizations are subject to external industry or government regulations for protecting sensitive data from exposure or loss. Companies in the card payment ecosystem, for example, must comply with the Payment Card Industry Data Security Standard (PCI DSS). Healthcare companies and organizations operating in the US must uphold the HIPAA Privacy and Security Rules. All publicly traded (and some privately held) companies with a US presence are subject to the Sarbanes-Oxley (SOX) Act, which requires measures to ensure financial data accuracy and safeguarding.

In addition to external compliance requirements, many companies establish internal governance policies to protect business-critical assets. E-commerce companies, for example, may adopt policies that prioritize vulnerability remediation for customer-facing assets.

For smaller organizations, with their limited vulnerability remediation resources, policy-driven campaigns are often the most appropriate because they are rule-based, repetitive, and lend themselves well to automation.

The main challenges of a policy-driven remediation campaign strategy are the need for continuous auditing and maintenance, as well as strict adherence to external and/or internal compliance policies. The following steps can help both large and small organizations mitigate these challenges:

- Map external and internal policy requirements to the company's existing vulnerability management program. Does your vulnerability scanning cadence and coverage align with these requirements? Do your vulnerability response workflows meet the time-to-remediation and other SLAs? Can you verify and document that a vulnerability has been successfully remediated?
- Scope the affected assets. PCI compliance requires quarterly scanning of external assets and rapid remediation of all vulnerabilities with a CVSS score of 4.0 or above. However, even PCI differentiates between critical systems for which patches must always be up-to-date and "less critical" systems that should be patched "as soon as possible." Remember, however, to document such decisions, noting why a system was not patched and when it is due for review.
- Once the environment, requirements, and SLAs are known, start automating procedures. For example, use multiple streams of internal and external data and advanced analytics to automatically generate organization-specific vulnerability risk scores, create cross-team vulnerability remediation playbooks that automatically trigger remediation activities in the correct sequence, or automate vulnerability remediation reports to the greatest extent possible. In all cases, continuously assess and optimize the automated procedures.
- Be sure to scrub the environment continuously to identify old vulnerabilities that were not successfully remediated or new vulnerabilities that now need to be dealt with. Document any compensating controls applied to vulnerabilities that cannot be fixed in the required SLA.

Focused product-driven campaign

For large organizations that govern hundreds of thousands of assets and, by extension, tens and hundreds of millions of vulnerability instances, product-driven remediation campaigns may be the best strategy. Although updating an OS version or patching a database across an entire organization is a complex undertaking that can take several months, it remediates hundreds and sometimes even thousands of vulnerabilities on a very large group of assets.

The product-driven remediation campaign strategy faces two significant challenges: end-of-life products or products with high technical debt can be hard to patch, plus the remediation workflow has to be minimally disruptive to business activities.

Implementation guidelines for optimizing product-driven campaigns include:

- Continuously monitor for new releases and incremental updates and assess the criticality to your organization of the security issues they address.
- Identify where you can accept certain risks. For example, the operational risk of patching end-of-life products may outweigh the cyber risk of the identified vulnerability. However, these decisions need to be assessed very carefully and should be reviewed continuously.
- Create a test plan that checks whether the new patched version will be compatible with all related existing services and applications. You want to avoid having to roll back a patch because it disrupted a dependent service or application somewhere in the organization's business environment.
- Deploy gradually across the company's assets, leveraging downtime windows judiciously in order to minimize downtime and disruptions.

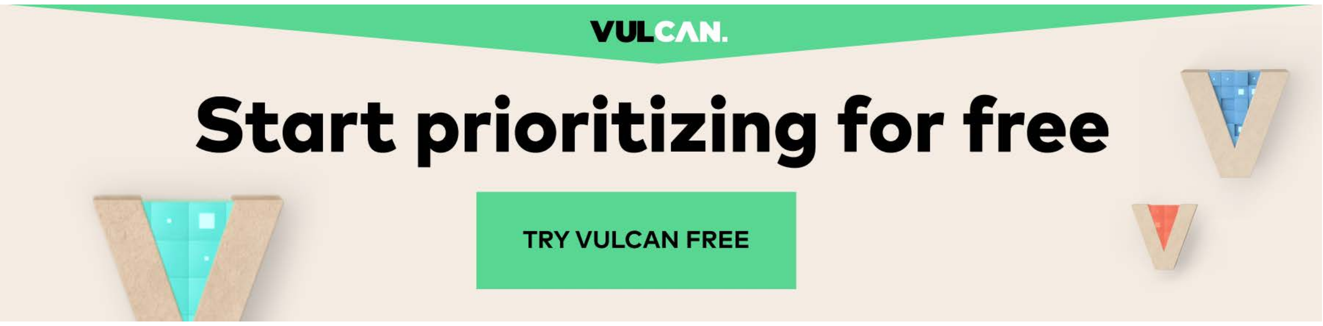

Choosing a remediation campaign

In an ideal world, an organization would remediate 100% of its known vulnerability instances in short time frames. In reality, we learn from Edgescan's 2019 Vulnerability Stats Report that, on average, it takes 77.5 days to close a vulnerability in the application layer (69 days for critical risks) and 81.75 days to close a vulnerability in the infrastructure layer (65 days for critical risks).

The key to effective vulnerability remediation is to focus on the vulnerability types that pose a great risk: active and dangerous exploits, threats that undermine the organization's compliance posture, and vulnerabilities that affect a great number of assets. One way to achieve that focus is to choose the most appropriate remediation campaign strategy by asking questions such as:

- How complex is your organization?
- How many different teams are involved in your vulnerability remediation processes?
- How many assets and vulnerabilities do you have to govern?
- How bound are you to compliance policies?
- How important is it to your organization to reduce its risk score?

Remediation campaigns with Vulcan Cyber

Vulcan Cyber orchestrates the entire cyber risk management lifecycle - from detection to intelligent risk assessment and automated resolution and reporting —for agile and targeted remediation campaigns. Whether dealing with urgent and high-risk vulnerabilities, complying with external or internal requirements, or patching assets at scale, Vulcan Cyber ensures that your vulnerability remediation campaign will always be targeted, optimized, and documented.

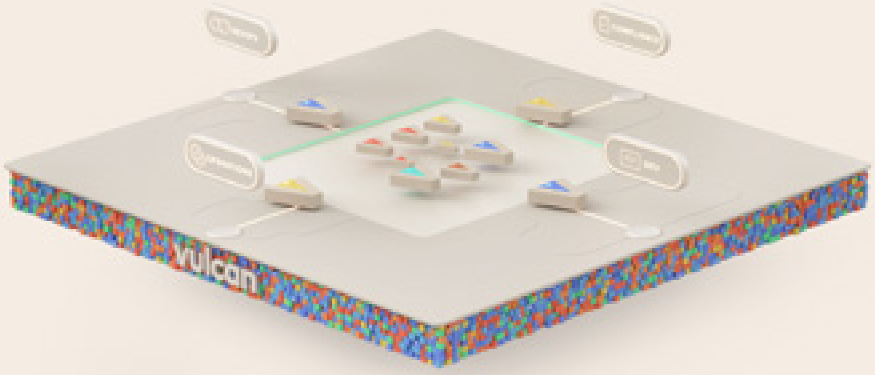

VULCAN