VERACODE
The Most Powerful Application Security Platform on the Planet

WHITE PAPER

Five Steps to Prepare for a Vulnerability Disclosure

Table of Contents

Executive Summary

The past year saw an increase in the number of branded vulnerabilities, with mainstream media covering stories of widespread software flaws. As a result, IT and security teams' responses to these vulnerability disclosures are now under more intense scrutiny from internal stakeholders as well as customers.

When a high-profile, or simply branded, vulnerability is disclosed, the security and risk teams are expected to temporarily abandon their planned activities and react. But, these responses can be time-consuming and costly. Despite the fact that the vulnerability may be achieving mainstream awareness, enterprises need to balance responses against risk.

This document provides guidance on preparing for a high-profile vulnerability disclosure so risk-management or security teams can respond with the appropriate level of urgency. Teams can use it as a starting point to formulate a strategy for vulnerability responses in the future so that they are prepared for the eventual disclosure.

Introduction: Preparing for a Vulnerability Disclosure

In 2014, vulnerability disclosures like Heartbleed, ShellShock and POODLE forced IT and risk managers to abandon planned activities so they could quickly patch their systems. In some cases, organizations had to first identify where the vulnerability existed within the company's infrastructure before they could start patching, costing more time and energy.

Early in 2015, several new branded vulnerabilities were disclosed. GHOST, FREAK and JetLeak were all given names and logos, but weren't as widespread, and none had a reliable widespread exploit available. IT and risk management teams were expected to react in real-time, regardless of the impact to the enterprise's planned activities. Without a comprehensive vulnerability response plan providing guidance on the appropriate level of urgency, IT and risk management teams were forced to treat all issues with a high level of urgency, just to be safe.

The number of high-profile, highly exploitable vulnerability disclosures shows no sign of waning, nor will the number of branded, yet less-urgent vulnerabilities. To ensure the appropriate level of urgency is subscribed to each vulnerability disclosure, IT and risk management teams need to create protocols for responding. Otherwise, the company risks having uncoordinated responses that grant the same urgency to low-priority issues as they would to highly critical vulnerabilities.

Veracode recommends taking a five-step approach to creating a vulnerability response program:

1. Identify a rapid-response team
2. Create protocols for the team to follow
3. Define priority levels and corresponding responses
4. Set up a framework for determining priority levels
5. Create clear procedures for responding to each level

Step 1: Identify a rapid-response team

Much like enterprises have incident-response teams in place to manage crisis situations such as data center outages or a security breach, enterprises should have two teams responsible for responding to vulnerabilities. The first-level, or rapid-response, team makes the initial decision about how to respond, while the second-level group, or security-incident response team, is the team of responders.

When a vulnerability is disclosed, the rapid-response team meets and runs through the process to determine how the company should respond and at what level of urgency. To ensure that as many viewpoints as possible are represented in the rapid-response team, the team should include a representative from each of the following functional groups:

- Security
- Engineering
- IT
- Finance leadership
- Customer facing leadership

The security-incident response team executes based on the level of urgency the rapid-response team subscribes to the vulnerability disclosure. This team typically includes members from:

- IT
- Security
- Engineering
- Marketing/communications

The IT, security and engineering members of the security response team focus on patching and mitigation, while the marketing/communication team member's main focus is to communicate the team's efforts internally and to customers when necessary.

Step 2: Create protocol for rapid-response team to operate under

In order for the rapid-response teams to operate effectively, each team should establish operating procedures. The rapid-response team's first responsibility is identifying when a new vulnerability is disclosed. This alert will typically come from a press release announcing the vulnerability disclosure or from a vendor letting you know what it is doing to respond. Either way, an internal email distribution list should be created, and any employee hearing about a new vulnerability disclosure should communicate this to the rapid-response team. Veracode suggests the following steps:

Step 1 – Alarm sounded by email to team distribution list.

Step 2 – Rapid-response team communicates to appropriate parties that they are investigating the issue and will provide guidance on how to respond to customers and prospects shortly. Also asks teams to not react until they hear back.

Step 3 – Rapid-response team reviews the situation based on all information available.

Step 4 – Rapid-response team determines what level response is required.

Step 5 – Rapid-response team communicates response plan to appropriate parties.

Step 3: Set up framework for determining priority levels

Every organization has its own appetite for risk, and, as such, each organization will have a different definition for what constitutes high urgency and what constitutes low urgency. Regardless of your organization's appetite for risk, a programmatic approach with clearly defined patterns will help in teams keep a level head throughout the response process.

See Appendix A for a list of questions that can be used to help determine what priority level a vulnerability falls into. See Appendix B for a flow chart template for helping to respond to the top-level questions.

Step 4: Define priority levels responses

Predetermining what constitutes a high-urgency vulnerability versus a low-urgency vulnerability will save time when an incident does occur. By explicitly outlining what actions will be taken given the priority level assigned to the vulnerability, enterprises can also justify their response to their customers and boards if asked.

See Appendix B for suggested priority-level definitions and their corresponding responses.

Step 5: Create clear procedures for responding to each level

Once the rapid-response team has determined what priority level the breach fits into, the security-incident response team will begin their mitigation efforts. It is important that there is a clear plan for how the team should respond depending on the level of urgency prescribed to the vulnerability. Veracode recommends the following responses:

Priority 1 Response:

A Priority 1 vulnerability requires an immediate response from internal teams. Given the high exposure, service providers and technology companies should also proactively communicate its response to customers.

Priority 2 Response:

A Priority 2 vulnerability requires a quick response, but may not require the enterprise to abandon planned activities. Service providers and technology companies should also proactively communicate its response to customers.

Priority 3 Response:

Patching efforts on a Priority 3 vulnerability can generally wait until the next scheduled maintenance activity. However, in the event customers inquire about patching activities, service providers and technology companies should be prepared with a statement on why it feels waiting is acceptable.

Priority 4 Response:

A Priority 4 vulnerability requires no response, as it does not impact the organization. A company would have a priority 4 response if they are not using the product that possesses the vulnerability. However, in the event customers inquire about patching activities, service providers and technology companies should be prepared with a statement on why the organization is not reacting.

For each response, the security-incident response team should follow these steps:

Step 1 – Determine which applications or systems need attention.

Step 2 – Prioritize patching efforts based on the applications' or systems' criticality to the business or likelihood of being exploited.

Step 3 – Communicate the response to all appropriate parties, which may include: employees, customers and other constituents that may be impacted by the response.

Next Steps

Once the enterprise has sufficiently answered these questions and determined what priority level the vulnerability falls under, it can begin responding to the vulnerability. To expedite the process, the organization should prepare detailed plans for how it will respond to each priority level in advance. These plans should include information such as which systems will be patched first, who will be on the response team, and how information will be communicated internally and externally.

After mitigation efforts are in place, there are still two questions the organization should consider.

Question 1: Am I tracking all my teams' mitigation/remediation activities?

If a vulnerability has existed long before the disclosure was made public, or became high profile, and the enterprise is working with an application security vendor, then there is a chance that many of the systems are already patched. Tracking patching efforts will help to determine if many systems are already patched and, thus, if an immediate response is required.

Question 2: Do I need to notify my customers?

This is a tricky question. If the enterprise wasn't breached, then the company most likely does not have an obligation to reach out to customers to inform them of its remediation attempts. However, if the vulnerability is high-profile and widespread, customers may begin asking for information about the enterprise's remediation and patching efforts. Enterprises should be prepared to tell customers how they are responding, and why this is the appropriate response. Having information about the exploitability of the vulnerability, as well as how many systems require patching, will go a long way toward making the case for why the enterprise responded the way it did.

Conclusion

Despite being a necessary part of security, responding to a zero-day vulnerability disclosure can be a costly effort if it pulls teams from their planned strategic or operational activities. Enterprises with pre-determined response plans and methods for prioritizing their remediation and patching efforts will be prepared to enact the appropriate response so that risk is properly mitigated with minimal cost.

Appendix A: Suggested Questions

These 10 questions will help enterprises to quickly determine the appropriate level of response and to develop a strategy for reacting to a vulnerability disclosure.

Question 1: What is the impact if the vulnerability is exploited?

For some vulnerabilities, the impact of an exploit is small to nonexistent. If the vulnerability is on a system that doesn't touch customers' or company data, then waiting for the company's regularly scheduled patch maintenance is usually a sufficient response to a disclosure.

Question 2: Is there already a public exploit available?

When a vulnerability is disclosed, it does not mean that this is the first time anyone has ever heard of it. The vulnerability disclosure will typically provide information regarding whether or not a public exploit is already available. If there is an exploit available, then responding to the vulnerability becomes a higher priority.

Question 3a: If so, how reliable or common is the exploit?

If an exploit is available, IT and security teams should find out how easy or common the exploit is. If it is not reliable or readily available, then the response doesn't have to be immediate.

Question 3b: If not, based on what is already known about the vulnerability, how soon can an exploit be expected to be created?

Just because an exploit isn't known, doesn't mean that the IT and security organizations shouldn't react. The team needs to determine how soon an exploit could be developed.

Question 4: Is it a protocol/design vulnerability (affects all implementations) or is it an implementation vulnerability (affects a specific number of implementations)?

Answering this question will help IT and security teams determine how widespread the vulnerability is in their environments.

Question 5: Can I reliably detect and/or block attack attempts?

In some cases, it may be sufficient to simply block attacks. However, this strategy is only effective in the short-term and should only be used to buy time for comprehensive patching and remediation efforts.

Question 6: Can I determine if the vulnerability has been exploited in my environment already?

If the company has already been breached by means of the recently disclosed vulnerability, then there are an entirely different set of questions and actions the teams must take[1]. Also, knowing that the vulnerability has been exploited should push the patching efforts up the priority chain.

Question 7: How many of my systems/applications are affected?

This is an important question to answer because it provides a scope for how big the patching efforts will be. Combined with the information obtained by asking the previous questions, knowing how big a job it is to patch all the necessary systems will help determine if the vulnerability disclosure requires an immediate response or can be scheduled for a later date.

Question 8: How should I prioritize patching?

While all systems should be patched eventually, IT and security teams should identify which systems are most vulnerable, which are critical to patch (public-facing, access to customer data, etc.) and if there are any systems that can wait for normally scheduled patching.

Question 9: Does our organization already have services in place that can help detect the vulnerability?

Finding all the instances of a vulnerability, especially one that is widely distributed, can be challenging. IT and security teams that already have application security testing solutions in place will have an easier time responding to a vulnerability disclosure. An existing method for finding a vulnerability will help answer many of the previous questions.

Question 10a: If so, can the vendor help us prioritize remediation attempts?

Not all application security solutions will have the ability to find all instances of a vulnerability — for instance, if the solution can only test internally developed applications and not third-party applications and open source or third-party components. IT and security teams should understand the capabilities and limitations of their application security solutions and work with their current vendor to determine what can be done to help respond to a vulnerability disclosure.

Question 10b: If not, what services are available/are there any special offers available to help respond to the vulnerability disclosure?

If the enterprise either doesn't currently work with an application security provider, or works with a vendor that cannot assess the security of third-party applications or components from open source or third-party libraries, then the teams should begin searching for alternative vendors that may be able to help them. In the case of a high-profile, branded vulnerability disclosure, it is likely an application security vendor that can assist in remediation activities will have a special offer to help its customers as well as enterprises it isn't currently working with.

Appendix B: Suggested Priority Definitions

Priority 1:

- Widely distributed vulnerability
- A reliable exploit is available or is expected to be available in the near future
- It is a protocol design vulnerability
- Impacts many systems
- Cannot reliably detect or block attacks

Priority 2:

- Widely distributed vulnerability
- A reliable exploit is not available or isn't expected to be available in the near future
- Implementation vulnerability
- Impacts several systems
- Can reliably detect or block attacks

Priority 3:

- Not widely distributed vulnerability
- A reliable exploit is not available or isn't expected to be available in the near future
- Implementation vulnerability
- Impacts only a few systems
- Can reliably detect or block attacks

Priority 4:

- The company does not have applications that possess this vulnerability

Appendix C: Flow Chart Template

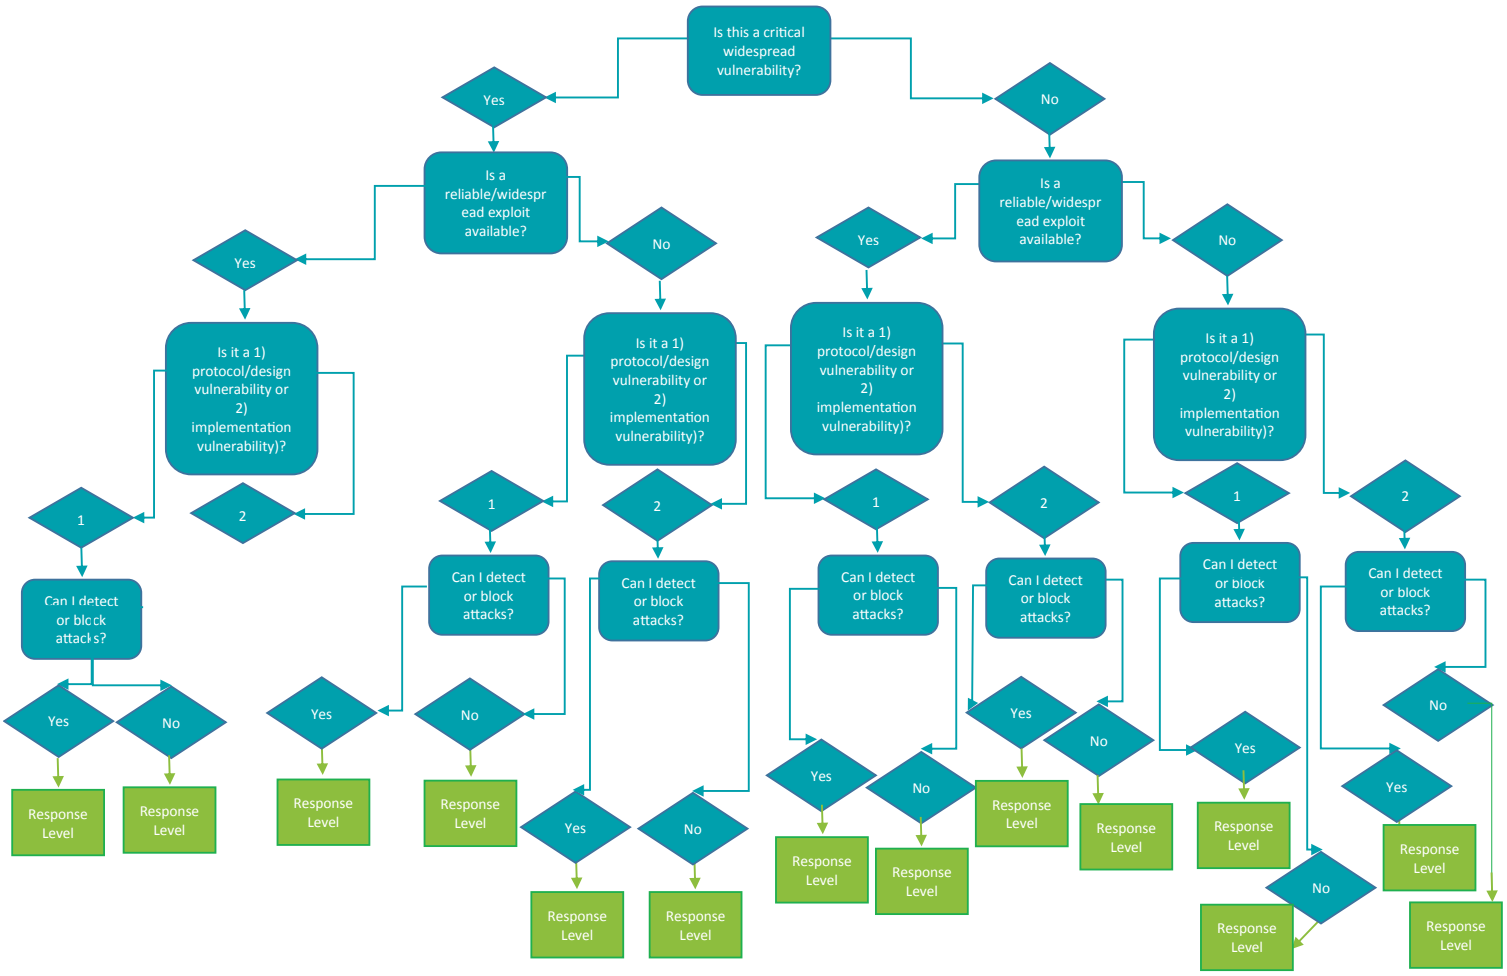

Customize this chart with the questions your organization feels are the most important, and assign the appropriate level response to each end point. Use this decision tree as a guide to determine what the enterprise's response should be.

End Notes

1. For more information on creating security plans, read Forrester's "Planning for Failure" Research Report. https://info.veracode.com/analyst-report-planning-for-failure-by-forrester.html