

Escaping the Doom Loop

Tim Willis

Head of Project Zero, Google

Charley Snyder

Head of Security Policy, Google

Eduardo Vela Nava

Head of Product Security Response, Google

Shailesh Saini

Head of Android Security Assurance, Android Security, Google

Andrew Whalley

Head of Chrome Security, Google

April 2023

Google

Summary

Looking beyond 0days

Driving patch adoption

Holistic lifecycle management

Normalizing transparency

Vendors should disclose when their products are actively exploited

More transparency around patching metrics will diagnose whether current approaches are working

Smart Transparency

Supporting researchers

The importance of intent in legal frameworks

Against gatekeeping

Escaping the doom loop requires more strategic approaches

The industry needs to improve at performing root cause analyses

Focus on the fundamentals

Conclusion

Summary

At Google, we work on security challenges across the full spectrum of cyber attacks – from spam and other nuisances which affect *billions* of people, to sophisticated exploits developed by highly professional teams to target the world’s most high-risk users. We don’t have the luxury of focusing on one or the other – improving trust online requires that we build mitigations that protect all our users.

Too often, we see public debate around security fixate on high-end threats and zero-day vulnerabilities, and not enough focus on the underlying conditions that enable them. Project Zero, our vendor agnostic security research team that studies zero-day vulnerabilities in hardware and software systems, is focused on “making zero-day hard,” but we see a need to develop new approaches to make *all exploitation* more difficult. Doing so requires working with a broad set of stakeholders: industry, who develop the platforms and services that attackers seek to exploit; researchers, who not only find vulnerabilities but identify and drive mitigations that can close off entire avenues of attack; users, who unfortunately still bear too high of a burden of security; and governments, who create incentive structures that shape the behavior of all these other actors. When we look at the ecosystem, it is clear that there is important work still to do in partnership with these stakeholders. We see four areas for improvement:

- **Looking beyond zero-days:** While zero-days continue to pose serious risk to society, more focus is needed to drive down the impact of vulnerabilities that are already known. The industry tends to focus on patching zero-days, rather than staying current on security updates as a whole. This practice can leave users open to harm and potential known vulnerability exploitation.
- **Normalizing transparency:** Time and again, transparency about attacks and vulnerabilities has proven essential to protecting users. More transparency about exploitation and patching is needed to protect users, understand whether current defenses are working and ensure that the defenses of tomorrow will at least nullify the attacks of today.
- **Supporting researchers:** While great strides have been made in recognizing (and protecting) the contributions of researchers, this progress needs to be built upon. The U.S. Justice Department has clarified their charging policies to [recognize](#) the positive contributions of security researchers, and this approach should be spread internationally and at the state level.
- **Escaping the doom loop:** The endless cycle of vulnerability, followed by patch, followed by vulnerability, is exhausting defenders and users. More investment is needed

to drive fundamental advancements in software security and speed the vulnerability-to-patch rate to escape this cycle.

This paper covers our thoughts in all four areas. These aren't just issues we are pointing out – we are committed to addressing them. That's why we are announcing the following initiatives today:

- ***Hacking Policy Council:*** For the first time, we are seeing laws (both [passed](#) and [proposed](#)) requiring the private disclosure of vulnerabilities to governments under certain circumstances. It is important that we get these laws right. That's why we are pleased to be founding members of the [Hacking Policy Council](#), a group of like-minded organizations and leaders who will engage in focused advocacy to ensure new policies and regulations support best practices for vulnerability management and disclosure, and do not undermine our users' security.
- ***Security Research Legal Defense Fund:*** Independent security researchers make enormous contributions to security, including at [Google](#), so protecting their ability to do their work is critical. We are proud to provide the seed funding to stand up a new legal defense fund to protect good-faith security researchers. "Good faith security research" means accessing a computer solely for purposes of testing, investigation, or correction of a security flaw or vulnerability in a manner that avoids harm to individuals and the public. Unfortunately, these researchers often still face legal threats when their contributions are unwelcome or misunderstood. Such threats can ignore the individual's rights or misconstrue facts, creating a chilling effect on beneficial security research and vulnerability disclosure, especially for those without resources. The [Security Research Legal Defense Fund](#) aims to help fund legal representation for persons that face legal problems due to good faith security research and vulnerability disclosure in cases that would advance cybersecurity for the public interest.
- ***Exploitation transparency:*** From time to time, vendors will release a fix without disclosing that the vulnerability was being actively exploited. Greater transparency around exploitation helps the industry better understand attacker behavior, ultimately leading to better protections. We believe this transparency should become part of the industry's standard vulnerability disclosure policies. We have always prioritized transparency when our products are exploited, but starting today we will make this an explicit part of our [policy](#), committing to publicly disclose when we have evidence that vulnerabilities in any of our products have been exploited.

Looking beyond zero-days

The life of a vulnerability doesn't end when the vendor releases a fix. Over the years, industry and policymakers have grown too focused on zero-day vulnerabilities as a top source of insecurity in the ecosystem. We need to shift to a more holistic approach to managing vulnerabilities focused on patching and software lifecycle management.

Driving patch adoption

Zero-day vulnerabilities continue to pose serious risk to the digital ecosystem, but the average user faces far more risk from known vulnerabilities that have not been patched. Whether you're an end user, an OEM, or a software provider, effective and timely patch incorporation is essential to hardening your security posture. Ultimately, the platform needs to release a fix to affected vulnerable parties to limit attackers attempting exploitation. Greater focus should be placed on the way platforms make patches available to users, including frequency of patching; options and incentives for automated patching; whether standalone security fixes are offered (versus feature updates); or whether app updates can be decoupled from full system updates for mobile devices. Project Zero, a vendor agnostic security research team that sits within Google and studies zero-day vulnerabilities in hardware and software systems, has pioneered patch and disclosure timelines for this very reason - for the immediate safety of users.

Ease of patch adoption in enterprise is a particularly understudied area of friction. Following many attack campaigns exploiting known, unpatched vulnerabilities, organizations are often chided for not applying patches in a timely manner. While this may be true, we tend to overlook some of the difficulties in patching. The industry should invest in making testing and applying patches easier for customers. Greater analysis of patch trends can help here. For example, this week we published a [deep dive](#) on Google Kubernetes Engine patching trends for Google Cloud customers, generating new insights and recommendations on addressing friction points.

Holistic lifecycle management

A vulnerability disclosure policy is a starting point for many companies, but we believe more holistic policies to address product life cycles must become the norm. Products should come with policies about expected lifetime (including expiration dates) and support and notification models for downstream customers. For instance, the Android team ensures that downstream partners (such as OEMs) have clear guidance on the security support timelines for the core Android OS (how long they can expect to get security patches provided by Google) as well as the Linux kernel (utilizing support timelines for long term support versions). We carefully select these to ensure that partners have a guaranteed period of support (minimum of 3.5 years)

from the launch of a specific version of the Android OS. Pixel also makes their specific update cadence available for [users](#).

Normalizing transparency

Transparency has proven essential to protecting users from online threats. Greater scrutiny by thousands of eyes produces digital products and services that are more secure, reliable, and trustworthy. Vendor transparency about vulnerabilities allows the development of ecosystem-wide mitigations and a shared view of attack trends.

Vendors should disclose when their products are actively exploited

If a vendor discovers a vulnerability being actively exploited (i.e. used by attackers to cause harm to users or organizations), it is not enough to just fix the vulnerability. Vendors should make users, supply chain partners, and the community aware of the exploitation and notify victims in a timely manner through public disclosure and direct outreach where possible. Making users aware of exploitation is especially important and time-sensitive when there are mitigations users can explicitly take to protect themselves against the threat, and the disclosure itself does not give attackers a significant advantage over defenders with respect to further leveraging the vulnerability. Additional details of vulnerabilities and exploits should be shared to improve researcher knowledge and defenses, weighing the balance of transparency and defensive benefit against the risk to users who are yet to patch. This is something we've prioritized at Google for years, and we've made it an explicit part of our [vulnerability disclosure policy](#).

More transparency around patching metrics will diagnose whether current approaches are working

More transparency from platforms around patch adoption metrics for users will help industry and policymakers understand the scope of the challenge and whether the industry is truly improving in this area. In enterprise settings, this should also include data around the amount of testing required for a given patch and rates of patch failures. Ideally, transparency would also extend to governments as they balance offense vs. defense considerations. The U.S. [Vulnerability Equities Process](#) and the Australian Government's [Responsible Release Principles](#) represent a positive step forward, but more data on outcomes could help further its mission. Other countries should follow the U.S.'s lead here but everyone should also improve upon it, such as by sharing the number of vulnerabilities disclosed versus those withheld from disclosure, or sharing more information about exploitation trends in general.

Smart Transparency

While transparency is in our DNA, our first principle is protecting users. We share information to raise awareness of threats and vulnerabilities, but sometimes sharing can put users at risk and add noise to the system if not done thoughtfully. We have seen recent policy proposals that would force companies to over-report events (e.g., report activities that provide no public interest benefit, such as scanning activity against public websites), or require the private disclosure of vulnerabilities to governments before customers are notified and prior to the development of mitigations. In the past, we have seen [well-intentioned policies have the opposite effect](#) — new policies in this area must be evaluated against their impact on security.

Supporting researchers

Industry and government have come a long way in recognizing the important contributions of security researchers to protecting users, systems, and organizations, but there are still outliers. We continue to see problematic attempts to criminalize or silence helpful research activities, or modify global best practice for vulnerability disclosure, for instance by compelling researchers to disclose vulnerabilities to the government before the vendor of the affected product.

The importance of intent in legal frameworks

Intent is important in these activities: testing a service to find vulnerabilities to contribute to a vulnerability disclosure program is different from testing to find vulnerabilities to exploit users. Legal frameworks that do not acknowledge the difference between research for defensive purposes versus malicious activities risk significantly chilling the former, which has become an essential component of the ecosystem. The United States has taken the lead in clarifying that security research should be [supported](#), not prosecuted, and this approach should be replicated elsewhere.

Against gatekeeping

We believe anyone, regardless of background, should be able to contribute to vulnerability research. Ultimately, vulnerability reports are information; organizations should not limit their ability to receive useful information from the community. While reports should be treated cautiously by the recipient organization, and bug bounty payments must follow all relevant legal requirements (e.g. relating to sanctioned entities), we oppose any efforts to “gatekeep” who can participate in vulnerability disclosure programs (for instance, by disallowing people with criminal records).

Escaping the doom loop requires more strategic approaches

Security can seem hopeless and endless at times: Vulnerability followed by patch; threat followed by mitigation. Each new attack trend spurs new solutions in the cybersecurity product market – but nothing seems to get better. We believe the best path out of this cycle of insecurity is not by bolting on new tools, but by focusing on the fundamentals of secure software development, good patch hygiene, and designing for security and ease of patching from the start.

The industry needs to improve at performing root cause analyses

At Google, we strive to eliminate entire classes of threats and vulnerabilities. This starts by performing root cause analyses of existing vulnerabilities to address the underlying architectural issues that allow them to proliferate. Too often, we see vendors apply incomplete fixes for serious vulnerabilities, addressing the symptoms of the issue without also treating the cause. This frequently leads to patch bypasses and waves of exploitation. For example, 17 of the 40 (42.5%) zero-days exploited in the wild which Project Zero analyzed in 2022 were variants of previously known bugs. This issue comes down to either a) a failure to understand the root cause of a given flaw, or b) a failure to prioritize truly fixing it. Focusing on root cause analysis will enable industry, government, and end users to start rising above the exhausting hamster wheel of vulnerability responses.

Focus on the fundamentals

Policymaker and industry attention can at times be reactive, with emphasis on addressing threats and vulnerabilities as they arise, rather than ensuring products are secure to start with. Fundamental software security practices do not get the attention their importance merits. Efforts such as those by the U.K. [National Cyber Security Centre](#), the U.S. [National Institutes of Standards and Technology](#) (NIST) and the [Cybersecurity and Infrastructure Security Agency](#) (CISA) to define and share examples of security best practices across the software development life cycle are a helpful step in this direction, but these efforts must be built upon. For example, a developer can follow all of NIST's Secure Software Development guidelines without ever considering whether to write their program using a modern, memory-safe programming language. As another example, Software Bill of Materials (SBOM) are a good start to understanding systemic dependencies and identifying insecure components – but SBOMs in and of themselves do not improve security. SBOMs should be a natural output of more secure and audited software build systems, and frameworks must be put in place to analyze SBOMs at scale.

Public/private partnerships are needed to develop flexible approaches to guide secure software development that work for organizations of all sizes, and internationally. Within Google, we continually update guidelines for developers to address evolving attacker techniques, such as the importance of [applying secure-by-design principles](#) during all phases of the software development lifecycle, and our evolution towards Rust for [new connected devices](#).

Conclusion

Though this paper references many challenges in the ecosystem, there is cause for optimism. Efforts like those from CISA reflect a growing desire to mitigate risk from both known and previously unknown vulnerabilities, and prioritize software security principles. Major platform providers have significantly [accelerated](#) the rate at which they develop and deploy patches. We are confident that the commitments we are making today, combined with a focus on the areas laid out in this paper, can drive significant improvements in vulnerability management, making the ecosystem safer for all users and organizations.