The Ethical Case for Penetration Testing Without Full Disclosure

Format: Argumentative | Level: Undergraduate | ~750 words

The security researcher who discovers a critical vulnerability in widely deployed software faces a decision with no good options. Immediate public disclosure allows defenders and attackers to respond simultaneously, providing equal opportunity to both patch and exploit the vulnerability. Coordinated disclosure through the vendor delays public knowledge but risks the vendor's failure to patch in good faith, leaving users vulnerable for an indefinite period. This essay argues that neither model is adequate as a universal standard, and that the appropriate disclosure timeline should be calculated as a function of the estimated patch deployment rate across affected systems rather than fixed at an arbitrary number of days.

The Problem with Immediate Disclosure

Immediate full disclosure of a vulnerability, including technical details sufficient to reproduce the exploit, maximizes pressure on vendors to patch quickly but also maximizes the window during which malicious actors can exploit the vulnerability before a patch is available. The argument for immediate disclosure rests on the assumption that vendors respond faster to public embarrassment than to private notification, which has empirical support in cases where vendors ignored private reports for months or years. However, this benefit is purchased at the cost of exposing every unpatched system to exploitation during the period between disclosure and patch deployment. For vulnerabilities in widely deployed infrastructure software, that period can affect millions of systems.

The Problem with Coordinated Disclosure

Coordinated disclosure, where the researcher notifies the vendor privately and agrees to a fixed embargo period before public disclosure, is the dominant industry model. The standard embargo period established by major technology companies is 90 days, after which the researcher may disclose regardless of whether a patch exists. The weakness of this model is that patch deployment is not the same as patch availability. A vendor may release a patch within 90 days while the majority of affected systems remain unpatched for months afterward, because enterprise software update cycles are slow and organizational change management is complex. Disclosing technical vulnerability details 90 days after vendor notification, based on the reasoning that a fix is available, ignores the reality that availability and deployment are different events with different timescales.

A Framework Based on Deployment Rate

A more defensible standard would tie the disclosure timeline to the estimated patch deployment rate rather than to a fixed embargo period. If a patch for a critical vulnerability reaches 80 percent of affected systems within 60 days of release, the risk profile of full technical disclosure at day 60 is substantially different from the risk profile of disclosure when deployment remains at 20 percent. This requires researchers and vendors to share data about deployment rates, which is a significant coordination burden, but it produces a more principled outcome: disclosure timing based on actual risk reduction rather than on an arbitrary deadline.

The Counterargument: Researcher Accountability

Critics of this framework argue that placing disclosure decisions in the hands of individual researchers, even researchers with good intentions, creates inconsistency and potential for abuse. A researcher who disagrees with a vendor's remediation timeline may use a deployment-rate framework to justify indefinite non-disclosure, or conversely, may disclose prematurely based on a disputed estimate. This is a legitimate concern. The response is not to abandon deployment-rate reasoning but to institutionalize it: independent security disclosure boards, modeled on the established processes for medical device vulnerability disclosure, could evaluate timelines based on evidence and provide a structured process for resolving disputes between researchers and vendors.

Conclusion

The binary choice between immediate disclosure and coordinated disclosure with a fixed embargo understates the complexity of the problem. A principled framework for vulnerability disclosure must account for the gap between patch availability and patch deployment, the differential risk to different user populations, and the practical constraints on vendor remediation timelines. Withholding technical vulnerability details beyond the standard embargo period is sometimes ethically justified when doing so demonstrably reduces risk to unpatched systems, provided that a credible mechanism exists for independently evaluating that judgment. A framework is not the same as a rule, and the ethical argument for disclosure flexibility is not permission for indefinite non-disclosure; it is recognition that a fixed calendar does not adequately capture the security risk landscape.