

Bug Bounty Programs: Paying Attackers to Be on Your Side

Sat, Aug 15, 10:00 AM EDT · <https://scottslab.io/posts/bug-bounty-programs>




TL;DR

A bug bounty formalizes outside-researcher access in a legal, controlled framework, and it works by aligning incentives: a researcher who'd otherwise have no safe way to report (or a tempting reason to sell) gets paid to hand you the bug instead. Programs run fully managed (the platform validates and analyzes reports for you) or semi-managed (they hand off to you earlier). The payouts are real: Google paid \$112,500 for a Pixel flaw; the DoD's "Hack the Army 2.0" surfaced 146 vulnerabilities for over \$275K.


A bug bounty program is how you turn the internet's curiosity into a security asset instead of a threat. It gives external researchers a legal, controlled framework to probe your systems and report what they find: scope, rules, and a promise you won't sue them for looking. Without that framework, a well-meaning researcher who trips over a flaw has a genuinely bad set of options; the bounty gives them a safe, rewarded one.

The reason it works is incentive alignment. The same person who could sell a vulnerability on the black market, or just walk away and leave you exposed, now has a reason to bring it straight to you: money, recognition, and legal cover. You're not hoping attackers are nice. You're making "report it to the owner" the most attractive move on the board.



Bug Bounty Programs: Paying Attackers to Be on Your Side

Operationally, programs come in two shapes. A fully managed program leans on a platform to triage incoming reports, validate them, and deliver you complete analysis. You get vetted findings, not a firehose. A semi-managed program hands off to your own team earlier in the lifecycle, which is cheaper and gives you more control but puts the triage burden back on you. The right choice comes down to whether you have the internal capacity to separate real bugs from noise.



Bug Bounty Programs: Paying Attackers to Be on Your Side

And the numbers aren't token gestures. Google paid a researcher \$112,500 for a single Pixel phone vulnerability back in 2018. The US Department of Defense's second Army bug bounty, "Hack the Army 2.0," ran in late 2019 (results announced January 2020). 52 hackers found 146 valid vulnerabilities for more than \$275,000 in payouts, which is a rounding error against what any one of those bugs could have cost in a real breach. (The first Hack the Army, back in 2016, drew nearly 400 hackers and paid out around \$100,000 for 118 bugs.) That's the pitch in a sentence: bounties are cheap compared to incidents.