bugcrowd

DATA SHEET

Understanding Bug Bounty Scope

When is an open-scope program a good idea?

A *scope* is the defined set of targets that have been listed by an organization as assets that are to be tested as part of a particular bug bounty engagement. Things that are listed as "in-scope" are eligible for testing, and things that are "out of scope" are to not to be tested. Within the context of a bug bounty, what's in-scope is what researchers are incentivized to report (and are rewarded for), and what's out of scope is off-limits with no compensation is given for findings.

Scope Types

- **Limited Scope** only includes single or specific targets, such as "example.com."
- **Wide Scope** includes a wildcard to the in-scope targets, such as "*.example.com." This signifies that any subdomain of example.com is in-scope. For instance, part of that wildcard could include previously unmentioned or unexplored attack surfaces such as "staging-2019.example.com" or "admin.example.com," both of which could have rich opportunities for identifying vulnerabilities. Hackers are particularly good at finding and exploiting assets that have been forgotten. By including the wildcard, your organization increases the probability of identifying security risks across a much broader swath of your attack surface.
- **Open Scope** has no limits on what hackers can test. Open scopes generally look something like "any externally facing asset belonging to Example Org" – where nothing is excluded, so long as it belongs to the organization. Hackers are highly effective at identifying assets here – some may find and exploit an old marketing page for an event from a decade ago; they may find keys or sensitive information stored on GitHub or Pastebin; there may be remnants from mergers, acquisitions, and other artifacts. Running an open scope leverages the power of the whole crowd to find and identify any exposures your organization may have, and most of the time, there's a lot more than you realize.

Most bounty programs tend to systematically evolve over time – starting with a basic, limited scope (example.com), moving to a more expansive, limited scope (accounts.example.com, api.example.com), then to a wildcard (*.example.com), and finally to an open scope ("anything belonging to Example org"). Some programs complete this journey in months, while others take years, but the important thing is that ¯¢}´¸-} always evolving.

Why Open Scope?

Threat actors don't have to play by any set scope or rules. Unlike the hackers in a bounty program, an attacker isn't limited to testing one part of the attack surface‰ x©| ¯hey always find the path of least resistance – which is most commonly not through assets that receive the most testing. If the goal of a bug bounty is to secure your assets by finding issues before the threat actors can, then both sides need to be playing by the same rules. For that reason, an open scope program is not only useful, but usually necessary. Furthermore:

- On average, programs with wider scopes (at least a wildcard) get nearly 250% more findings than programs with limited scopes. This includes nearly 250% more P1s – where each and every P1 identified represents a breach that could have happened, but didn't.
- Programs with wider scopes often have longer and more substantial hacker engagement (more than 2x that of limited-scope programs), with testers submitting more findings over a longer period of time.

Finally, hackers are extremely good at finding hidden issues. Often, when organizations move to wide or open scopes, they learn about a significant number of assets that were simply forgotten or lost.

FAQs

What I don't want to invite hackers to inspect everything?

This is the most common concern, but the reality is, threat actors don't need your permission. The least you can do is to level the playing field, so that good actors (responsible hackers) can help secure these assets before they're exploited in the wild.

If I put all my assets in scope, won't I be neglecting the critical ones that I care the most about?

This is a reasonable concern; however, the simple answer is to make it more valuable to test on/against the things you really care about by tiering out the reward structure. If it pays many times more (say, 5x) to find bugs on the main app, testers will still probe around the rest of the attack surface, but they also know that the big money is where you put it. In this way, you can work to secure all your assets, but also emphasize the ones you care the most about.

I like the idea of having an open scope, but won't I run out of money?

This is also a reasonable concern. But given the option to pay $X to know about a critical vulnerability against a system that's not on your radar, or to deal with a multi-million dollar breach in the aftermath, wouldn't you prefer the former? Keep in mind that with a bug bounty you only pay for valid vulnerabilities. Meaning, if there's nothing out there, then there's nothing to reward... and if there is something, your rewards are set up in a way so that they align with the value that's being derived from the finding.

Why would I need open scope for a small attack surface?

Regardless of size, it can still be valuable for any organization to run an open scope program. If there's not much out there, then there's not really any downside to running open scope, so why not level the playing field and do it anyway?

Additionally, as you're likely well aware, in today's cloud-based, remote-work environment, an organization's attack surface is extremely complex and always evolving. There's a lot of exposure out there, and the crowd is the best kept secret weapon when it comes to helping organizations secure the totality of their footprint.

How do I start moving towards an open scope?

The best place to start is by talking to your Bugcrowd Success Team – your TCSM will help provide guidance, recommendations, and support for whatever you need to get going. All you really need to get started is an appetite for doing so – we'll help with the rest. As a note when opening up scope: it can be helpful to provide a list of assets you already know to be in scope. This gives researchers a starting point and saves them from having to do some early legwork. (It also allows them to focus on what really matters.) Generally speaking, the more information you can provide, the better.

Bugcrowd is happy to help champion it internally for you as well. If you need data, quotes, references, or anything else to help sell an open scope internally, we're happy to help however we can. We're here to help you secure your organization, and we truly believe that going to an open scope is a critical part of that security journey, and want to help however we can.