Crowdsourced Security — The Good, the Bad, and the Ugly

SOURCE Boston
April 26, 2017

Mike Shema
mike@cobalt.io

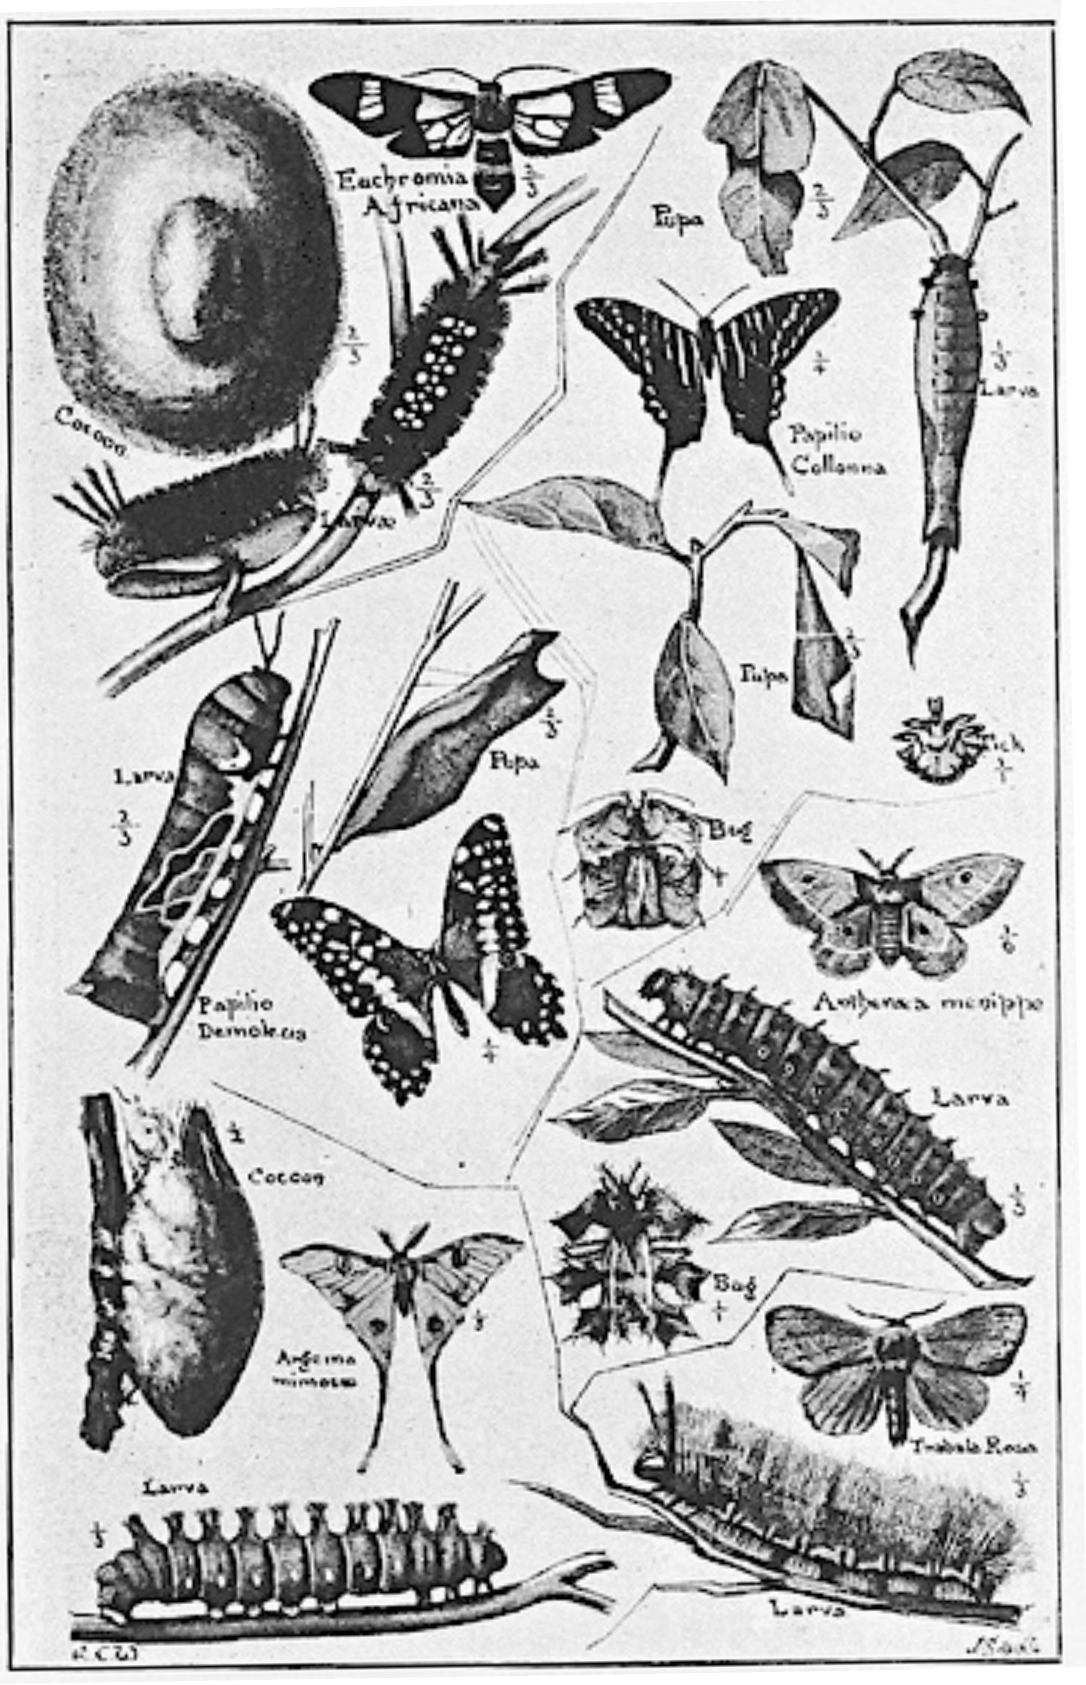

Bug Bounties embrace a crowdsourced model for discovering vulns.

They reward those who disclose vulns in a way that minimizes risk to the app, its data, and its users.

…and are just one of many alliances.

Uneasy Alliances

“What’s the price for this vuln?”
— Bounties

“What’s the cost to fix this vuln?”
— DevOps

“What’s the budget for finding vulns?”
— CISOs

“You see, in this world there's two kinds of people, my friend: Those with loaded guns and those who dig. You dig.”

– Clint Eastwood; The Good, the Bad, and the Ugly.

> “There are two kinds of spurs, my friend. Those that come in by the door; those that come in by the window.”
>
> – Eli Wallach; The Good, the Bad, and the Ugly.

How do we discover and fix vulns efficiently?

Noisy crowds produce a high rate of reports with a low percentage of vulns.

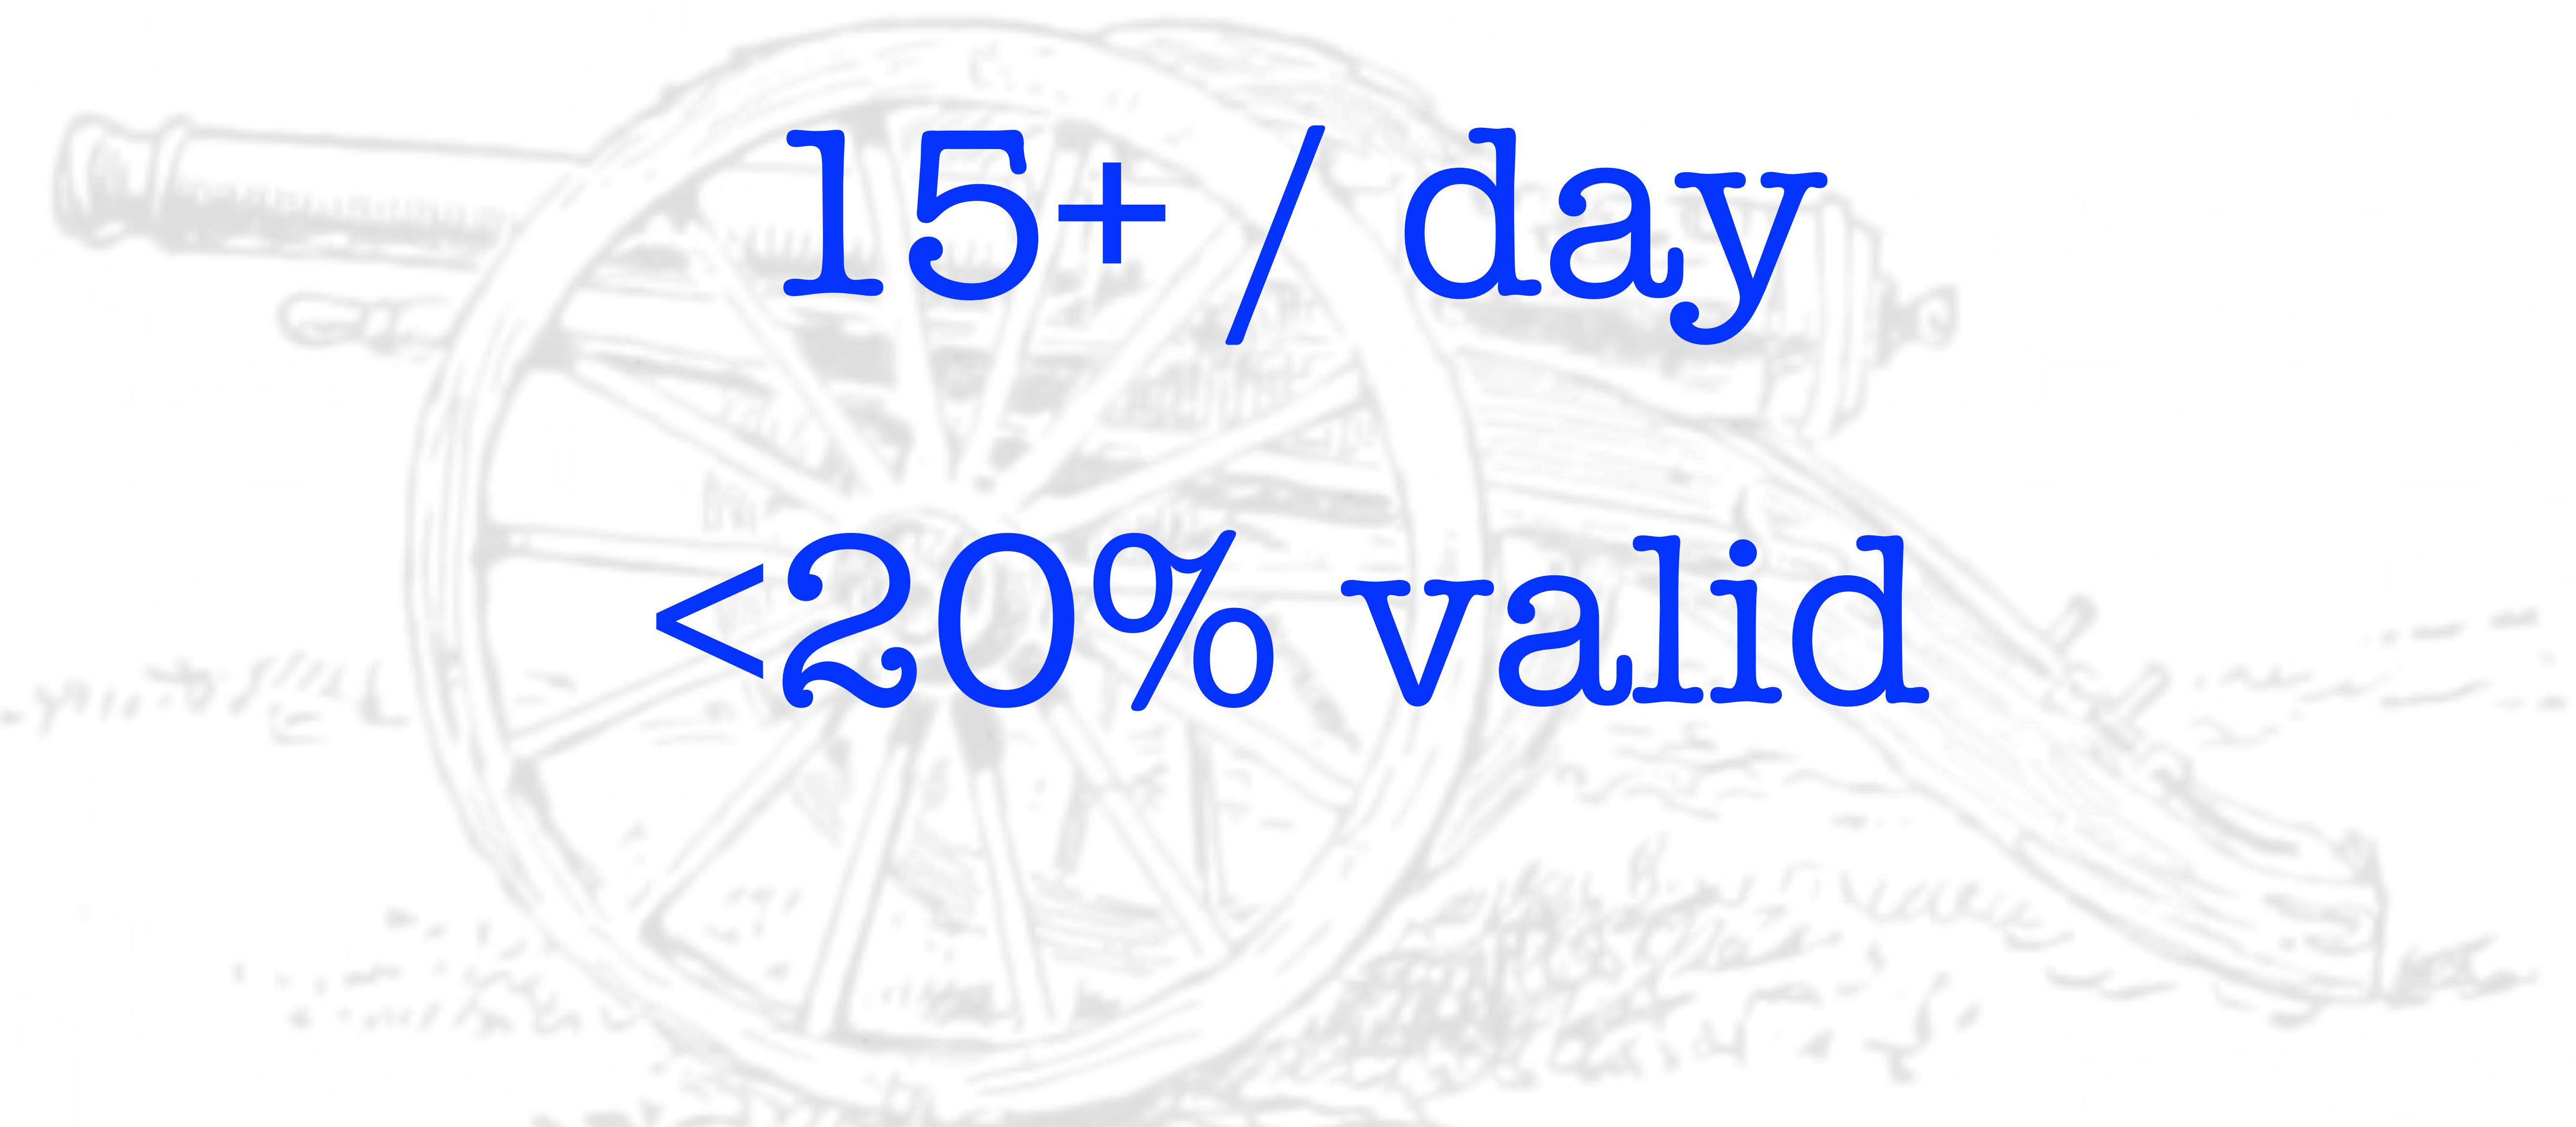

15+ / day

<20% valid

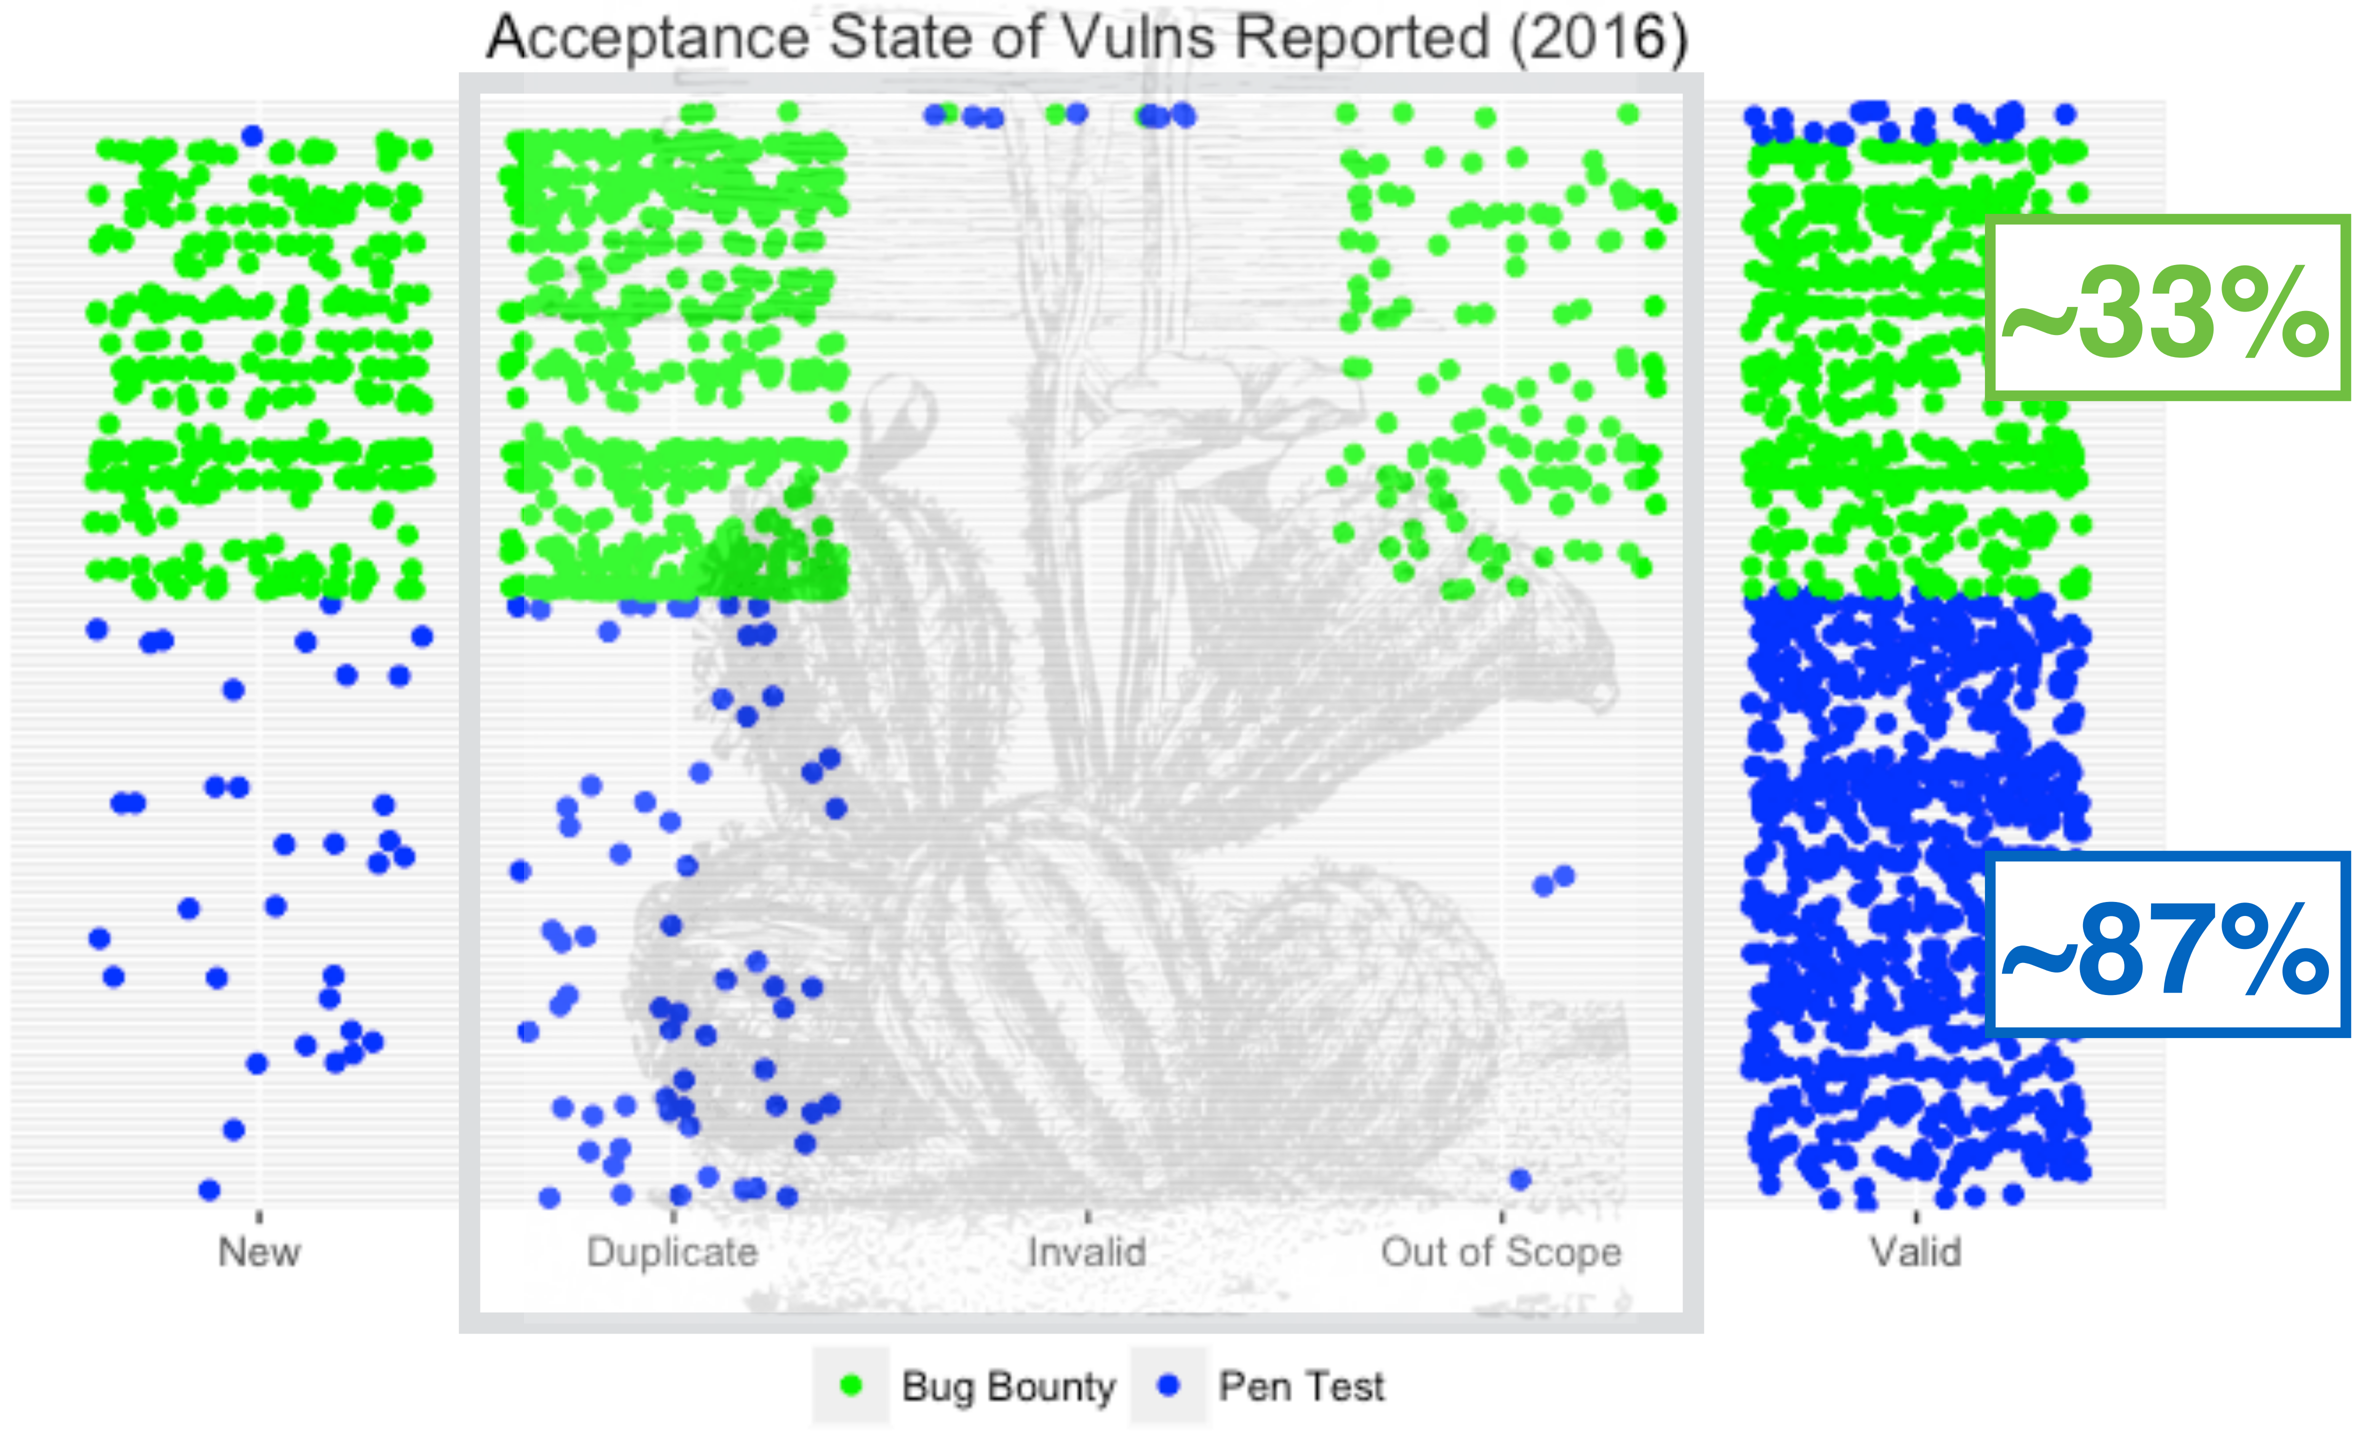

Reduce Noise

Create barriers to entry.

Filter crowds with selection criteria.

Bounty awards become a proxy for risk, where price conveys relative impact.

$0 — $15K

~$800 avg.

$50
Reflected XSS, self, no auth

$10,000
XSS vs. any auth’d user, access sensitive info

Bug Value

Pricing awards is one component of risk analysis.

Vulns reveal gaps where the SDL can be improved.

Metrics support mature SDL efforts.

Costs of Discovery

Bounties pay for the vuln found, not the effort to find it.

Discovery is haphazard.

What if we adjust the crowd?

Building Alliances

Diplomatic Missions

“This generates $X million in revenue. It’s not supported due to org changes.”

“It’s an internal system.”

“No one’s using it; we’ll just shut it down.”

“This was EOL last quarter.”

“This will be EOL next quarter.”

“We’ll accept the risk.”

Choosing Crowds

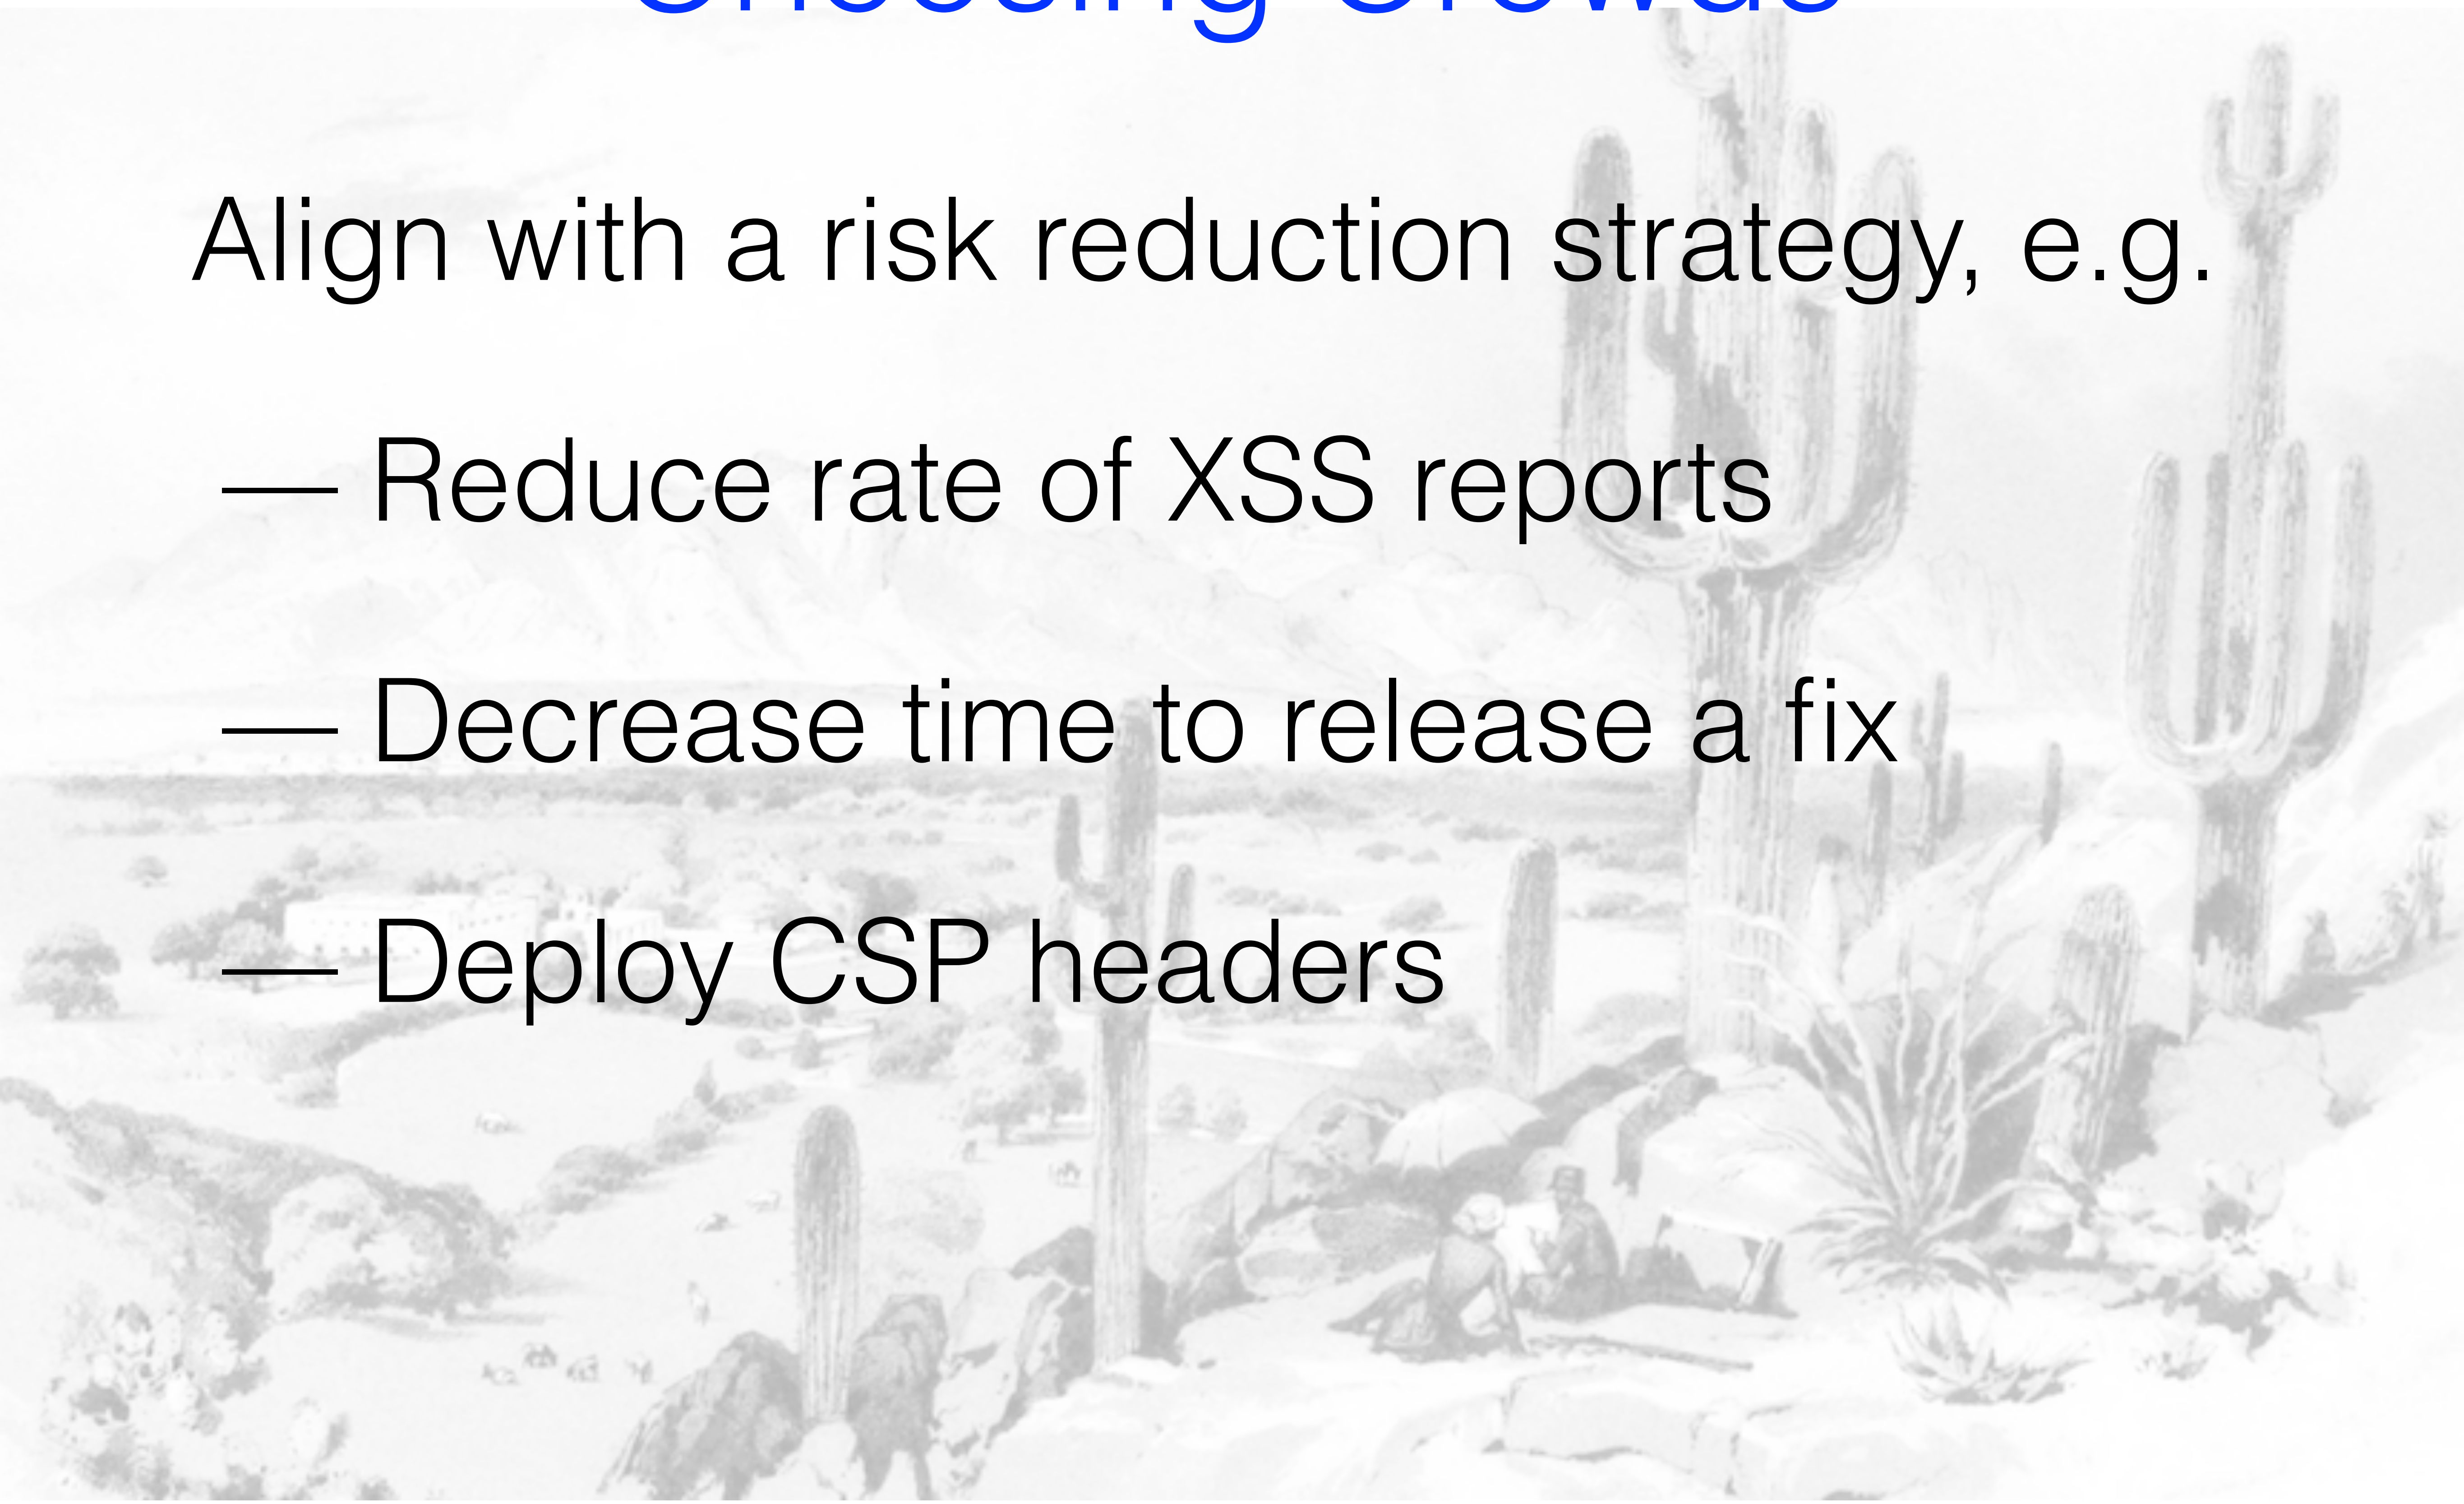

Align with a risk reduction strategy, e.g.

— Reduce rate of XSS reports

— Decrease time to release a fix

— Deploy CSP headers

Measuring Crowds

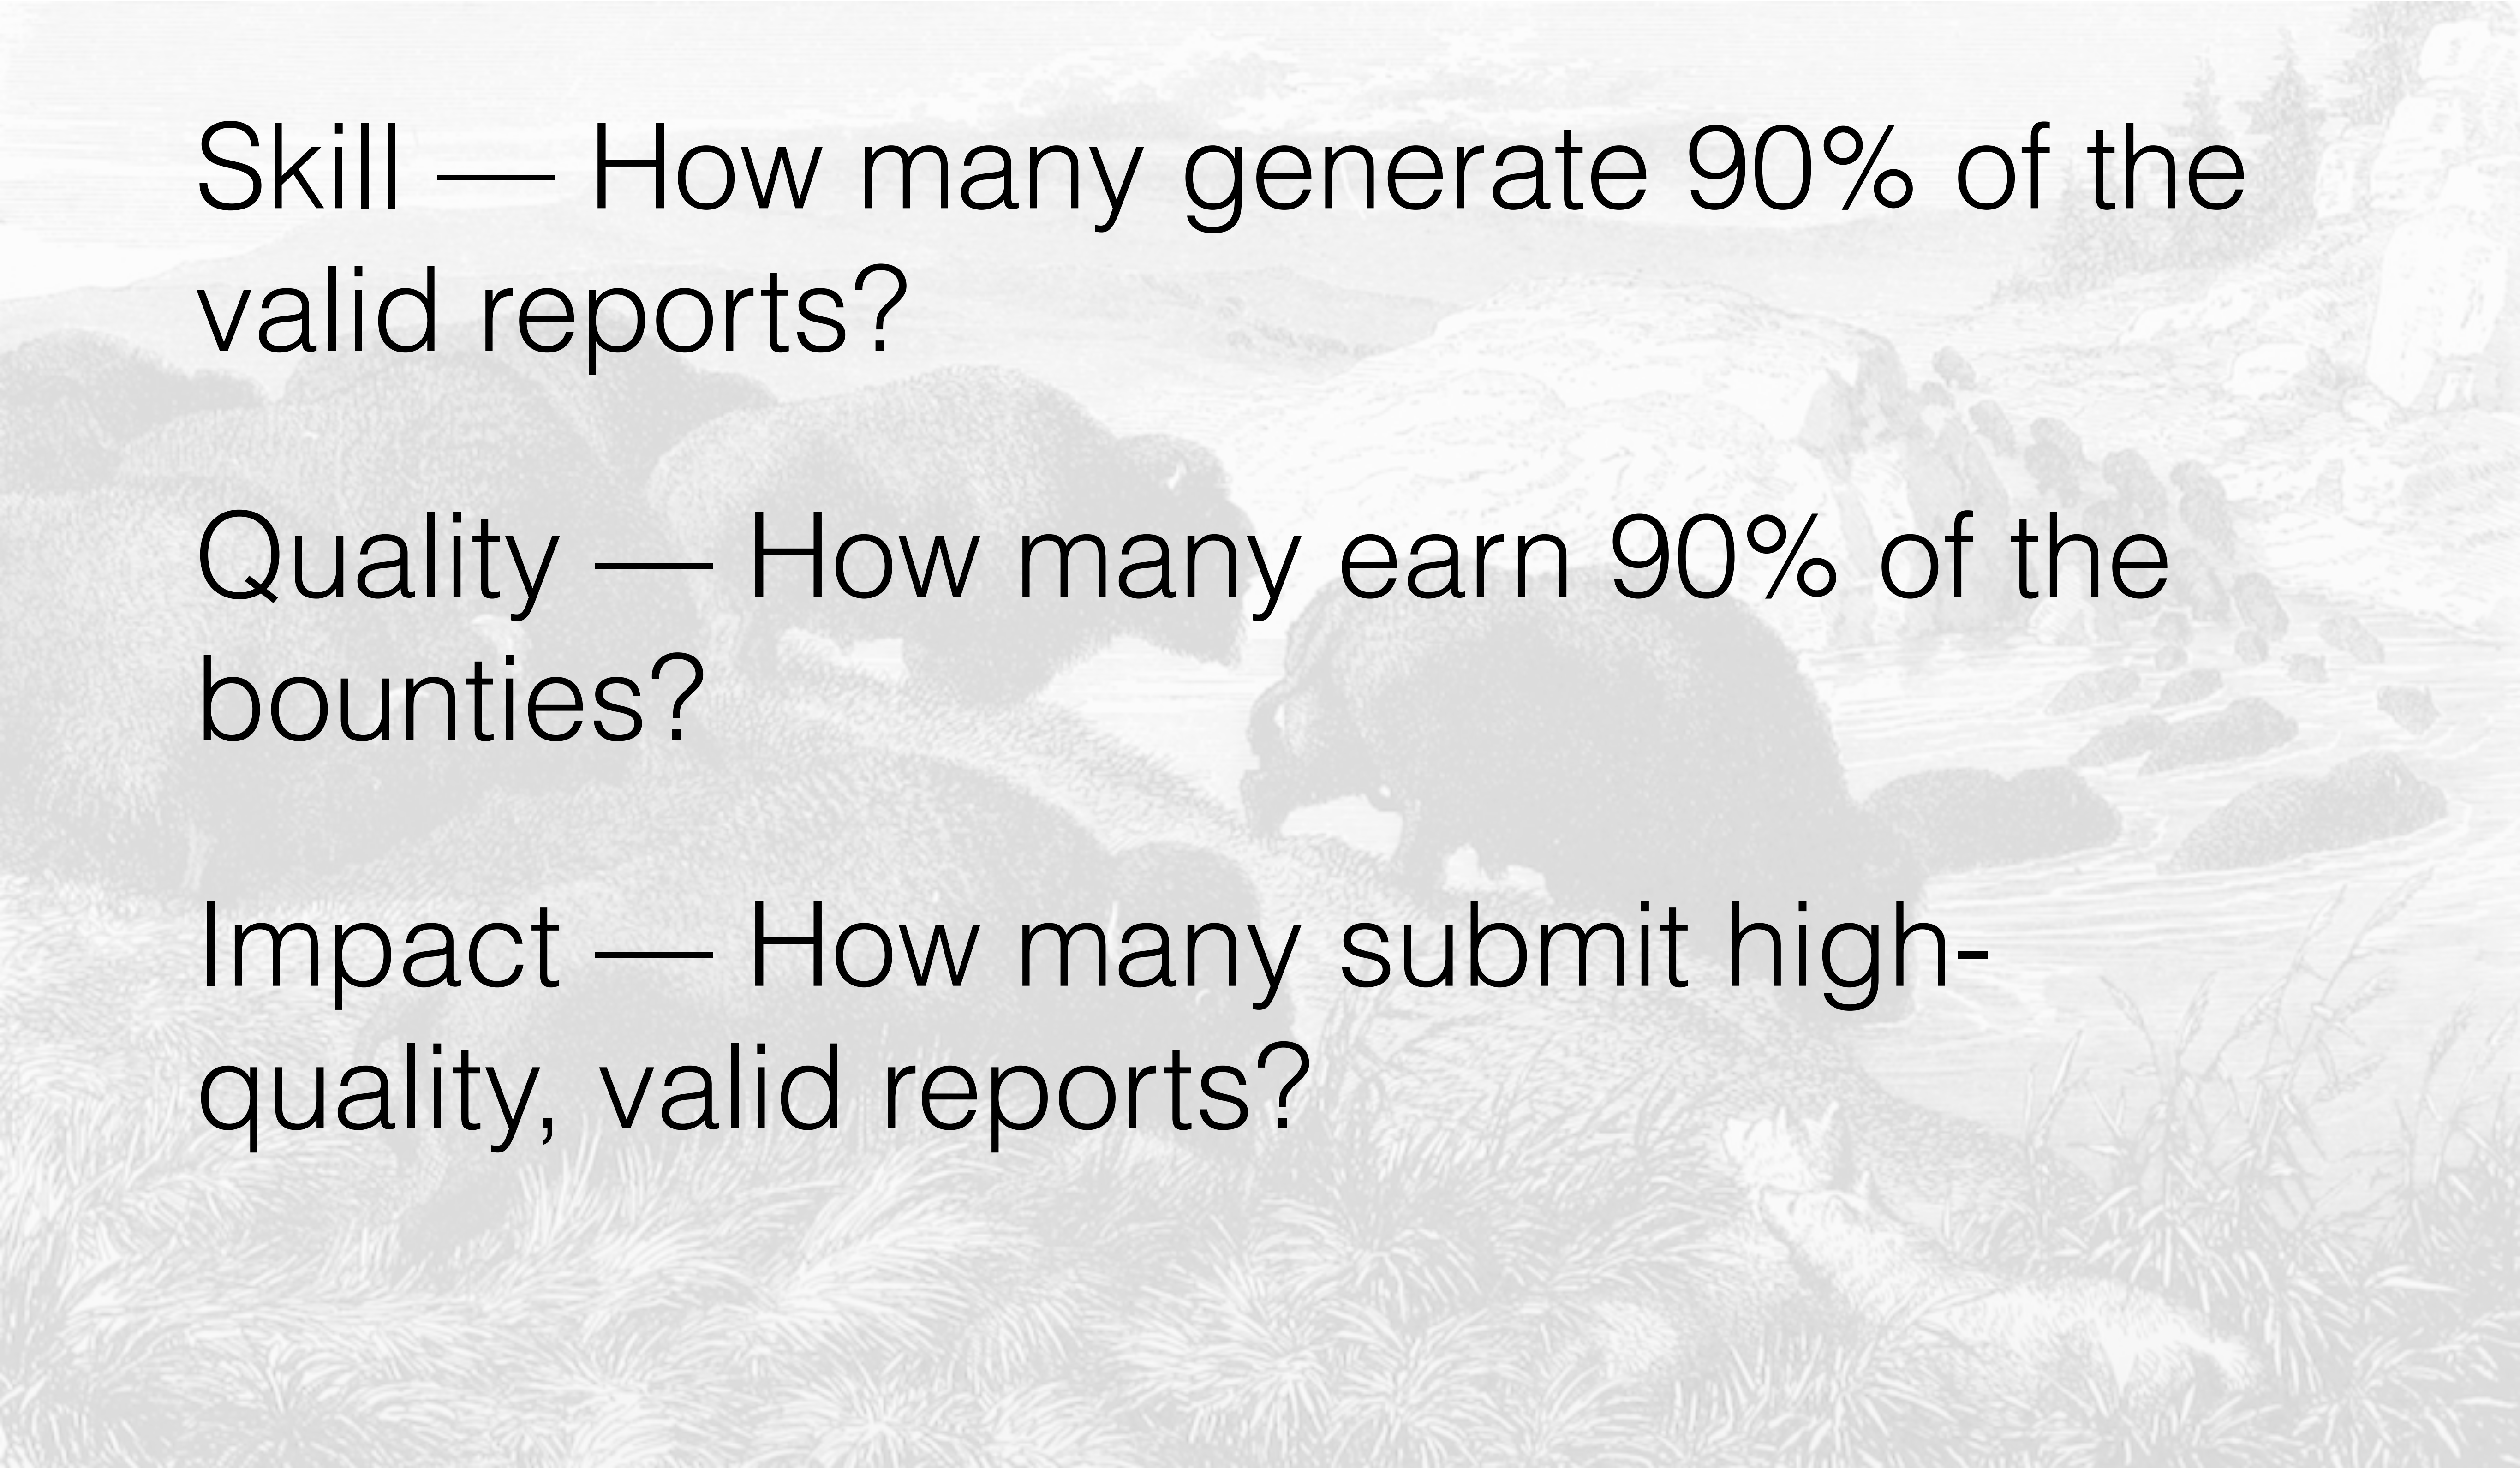

Skill — How many generate 90% of the valid reports?

Quality — How many earn 90% of the bounties?

Impact — How many submit high-quality, valid reports?

Crowd Containers

Size

Knowledge

Skill

Productivity

Documentation

Quality

Time

Compensation

…

Crowds that know ____ technologies, have ____ skills, to perform ____, and work for ____ incentives.

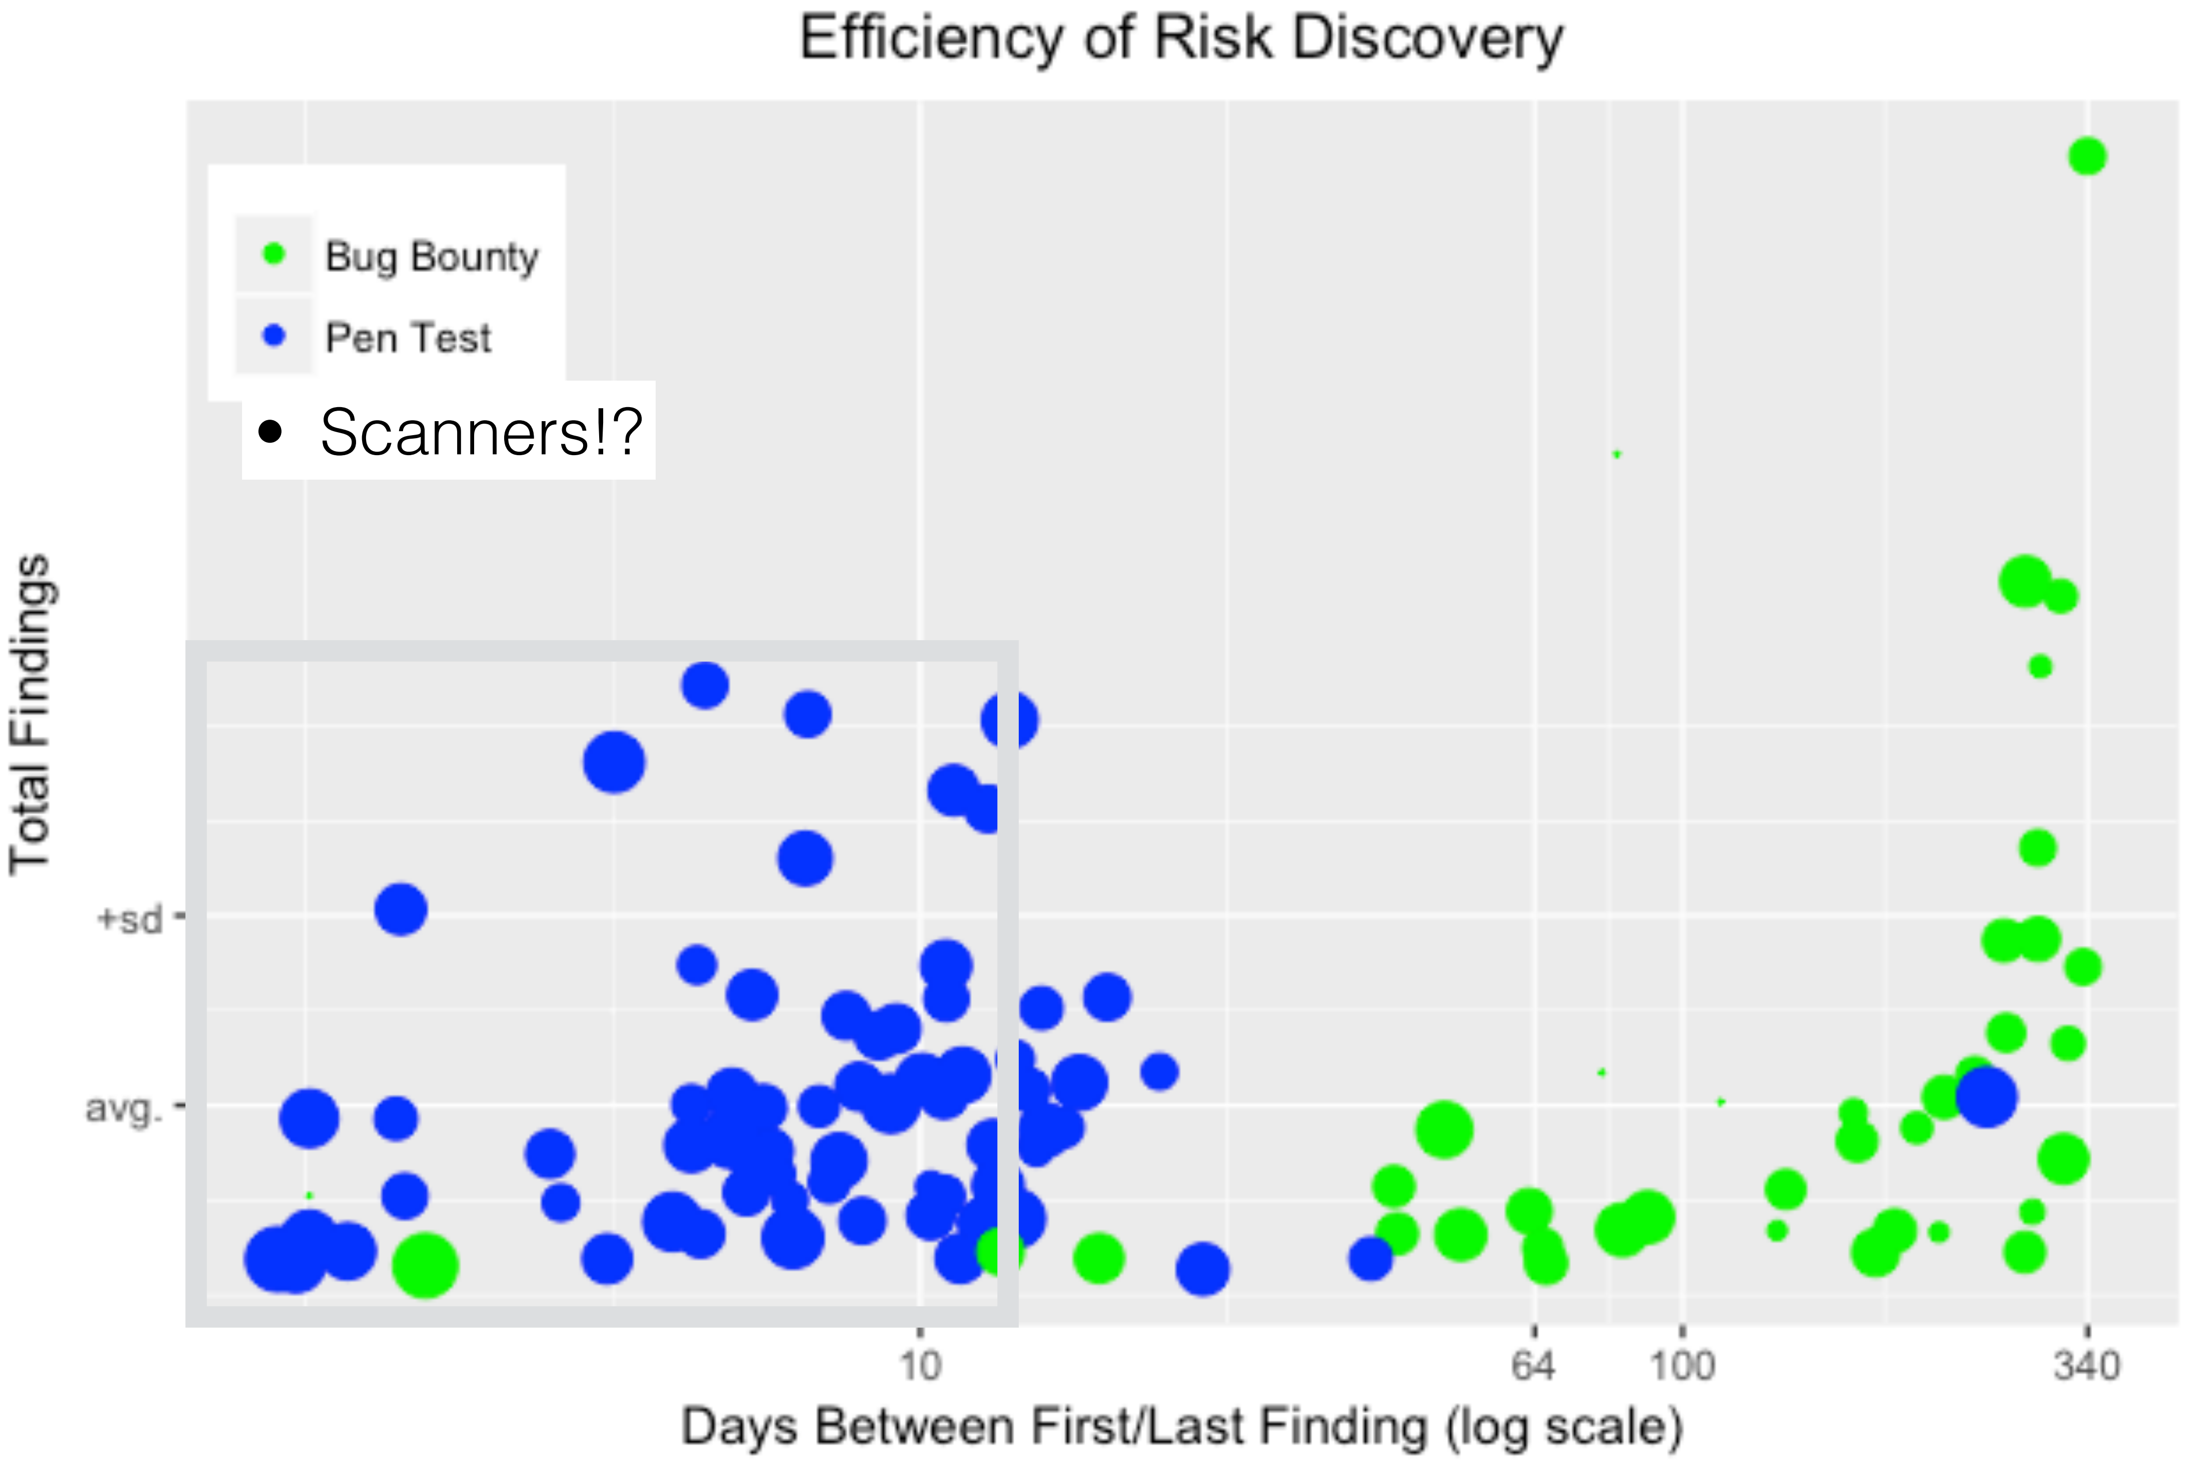

Find Faster

Working with release cycles.

Integrating with DevOps tools.

Exhausting the Pace of Vulns...or Attention?

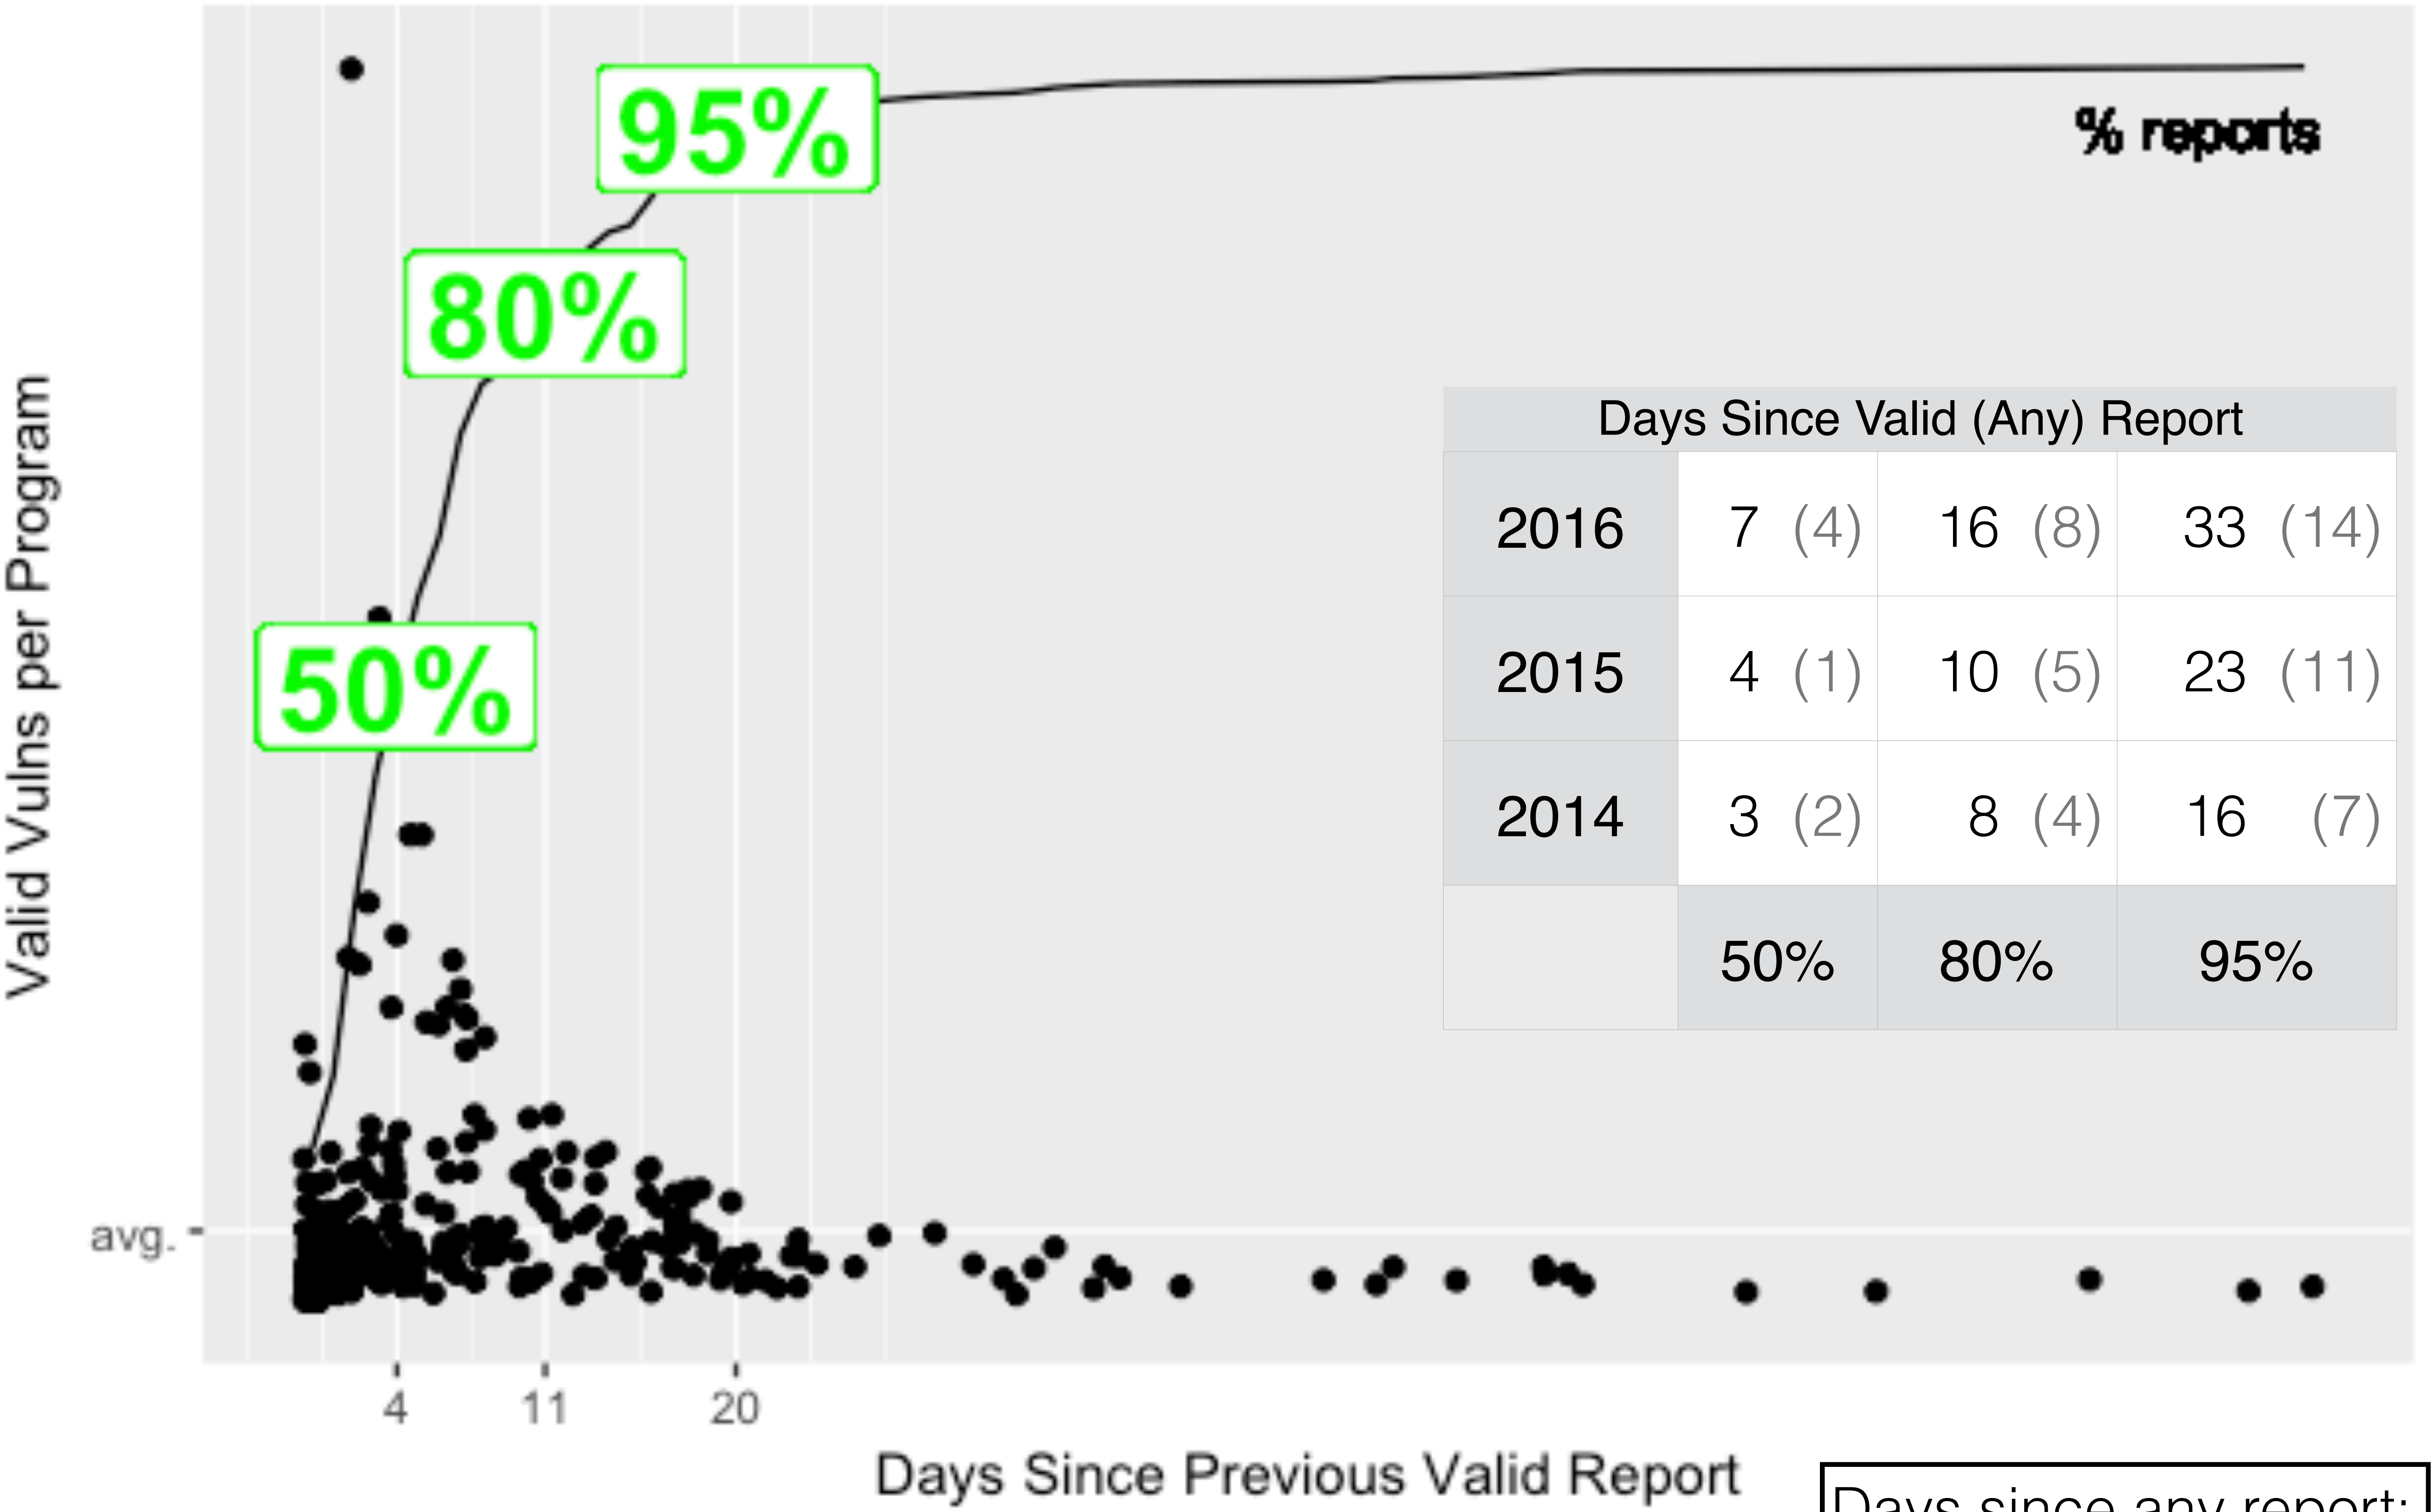

Days Since Valid (Any) Report			
2016	7 (4)	16 (8)	33 (14)
2015	4 (1)	10 (5)	23 (11)
2014	3 (2)	8 (4)	16 (7)
	50%	80%	95%

Days since any report:
2, 5, 11

Endemic Risk
(aka Mike's Tragic Quadrant)

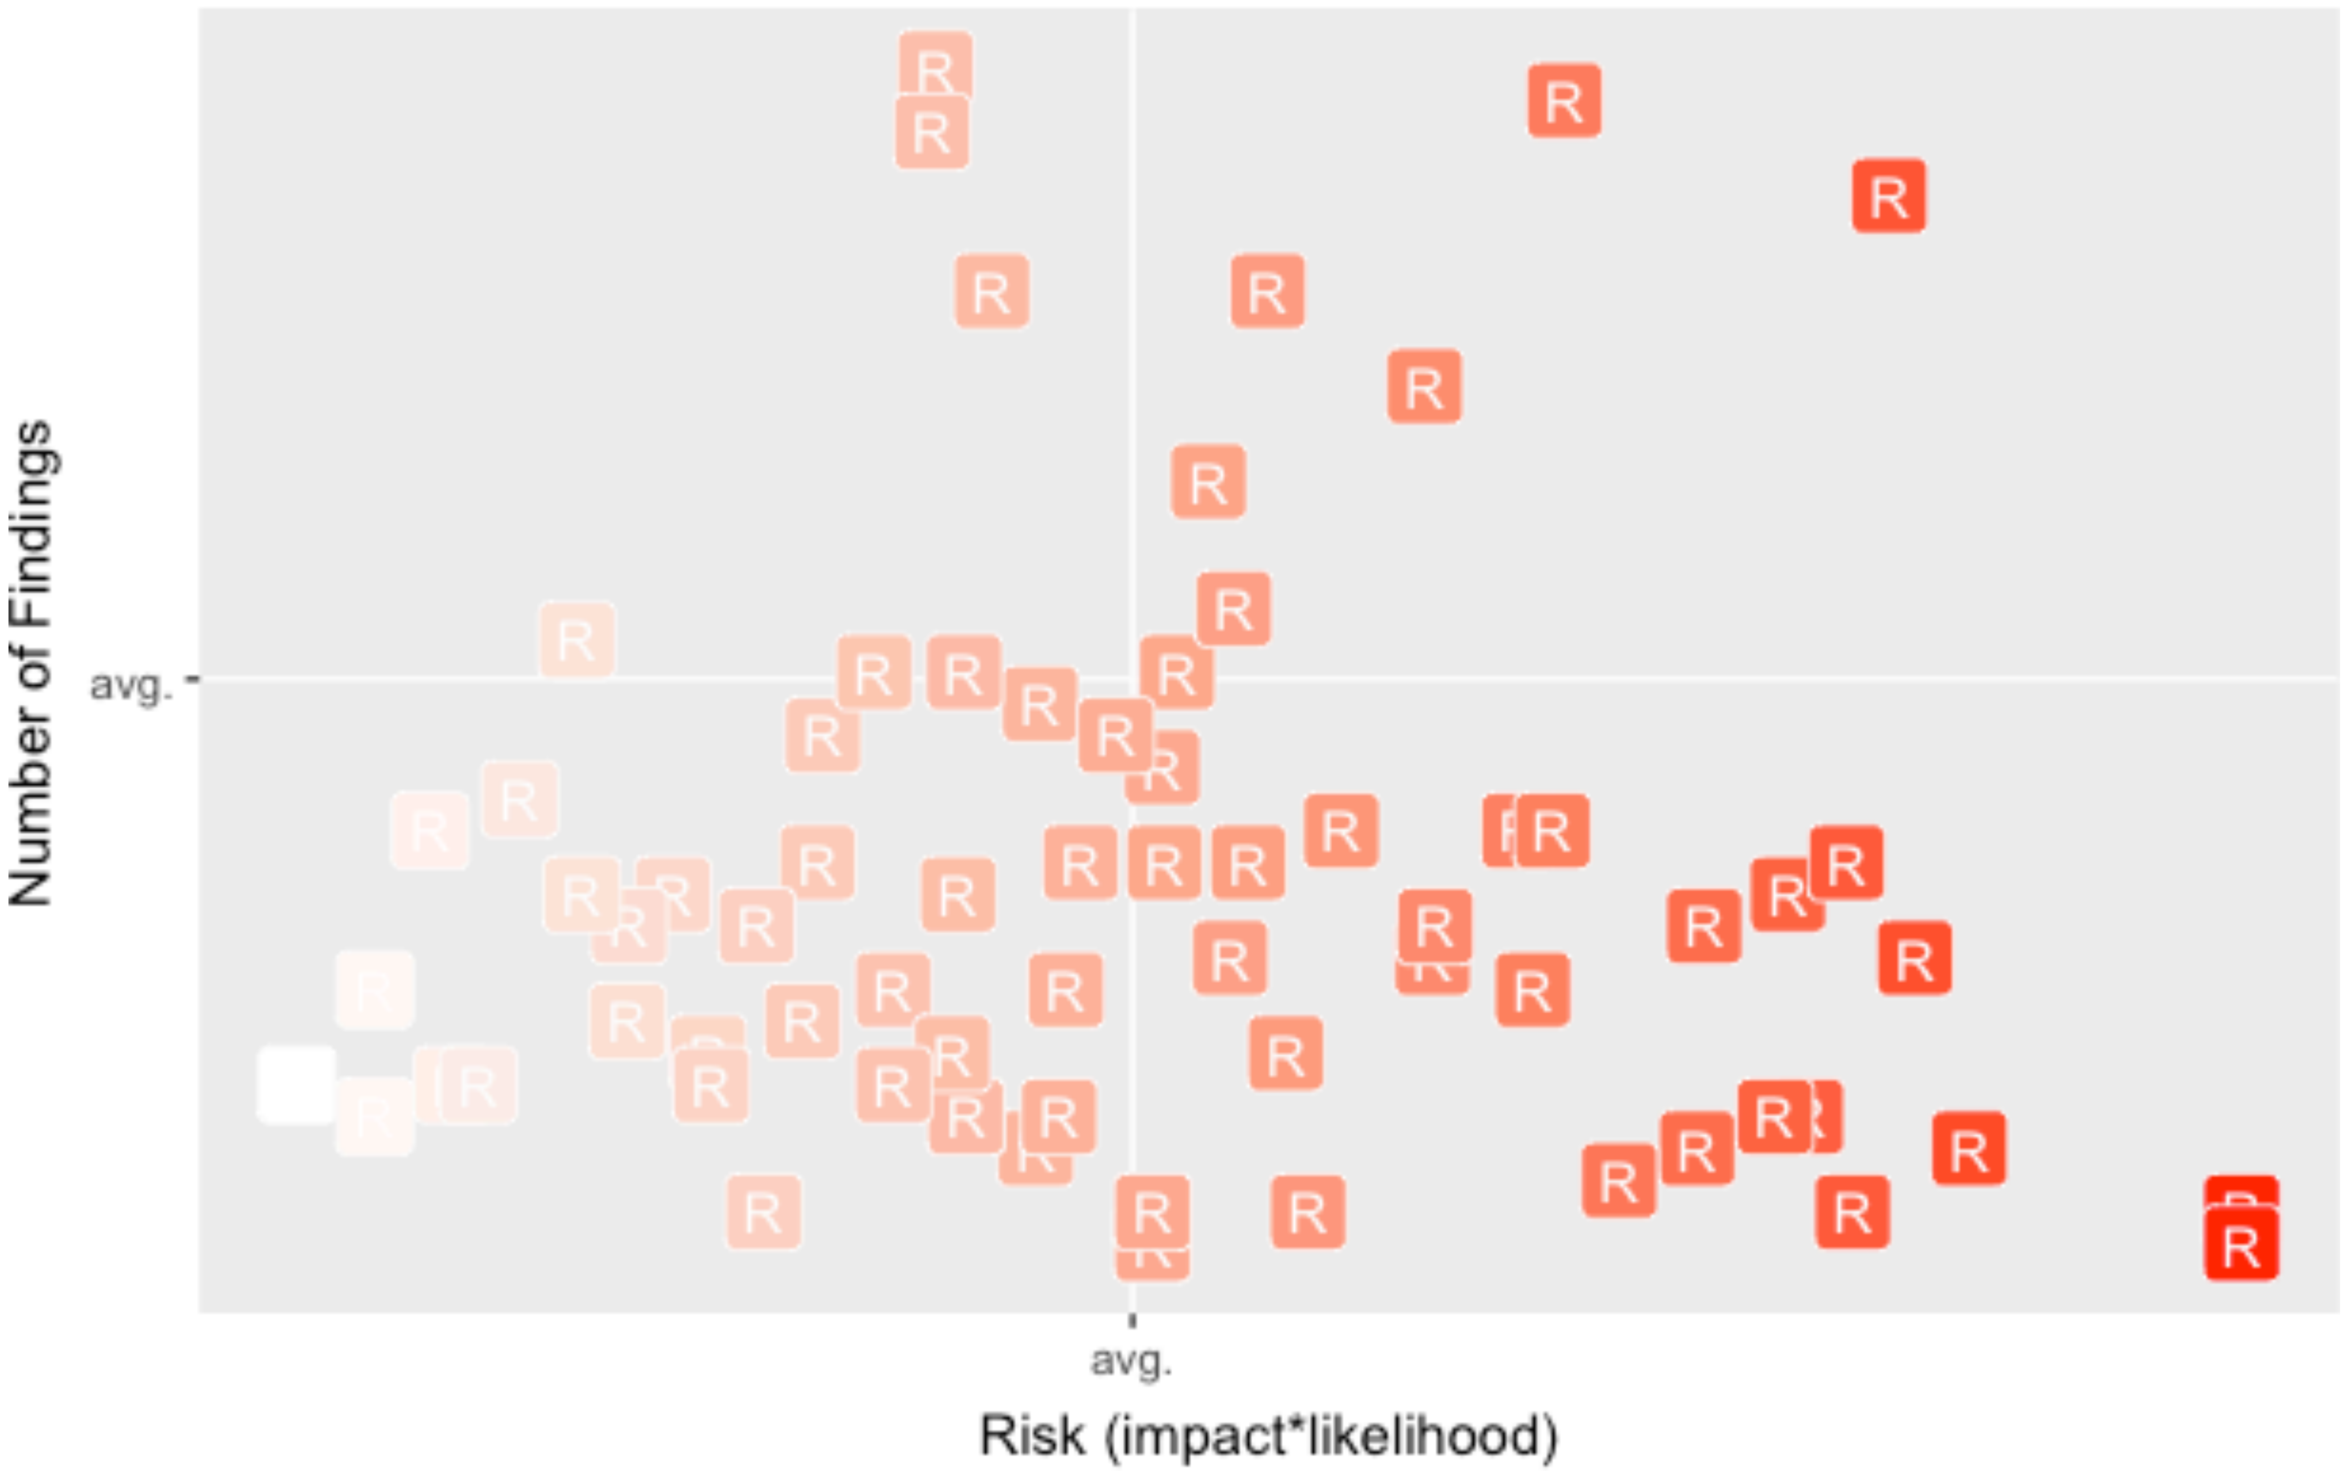

Track the sources of high-quality, high-risk reports.

Track the speed of deploying fixes.

Track trends in the amount and average risk of vulns in your app.

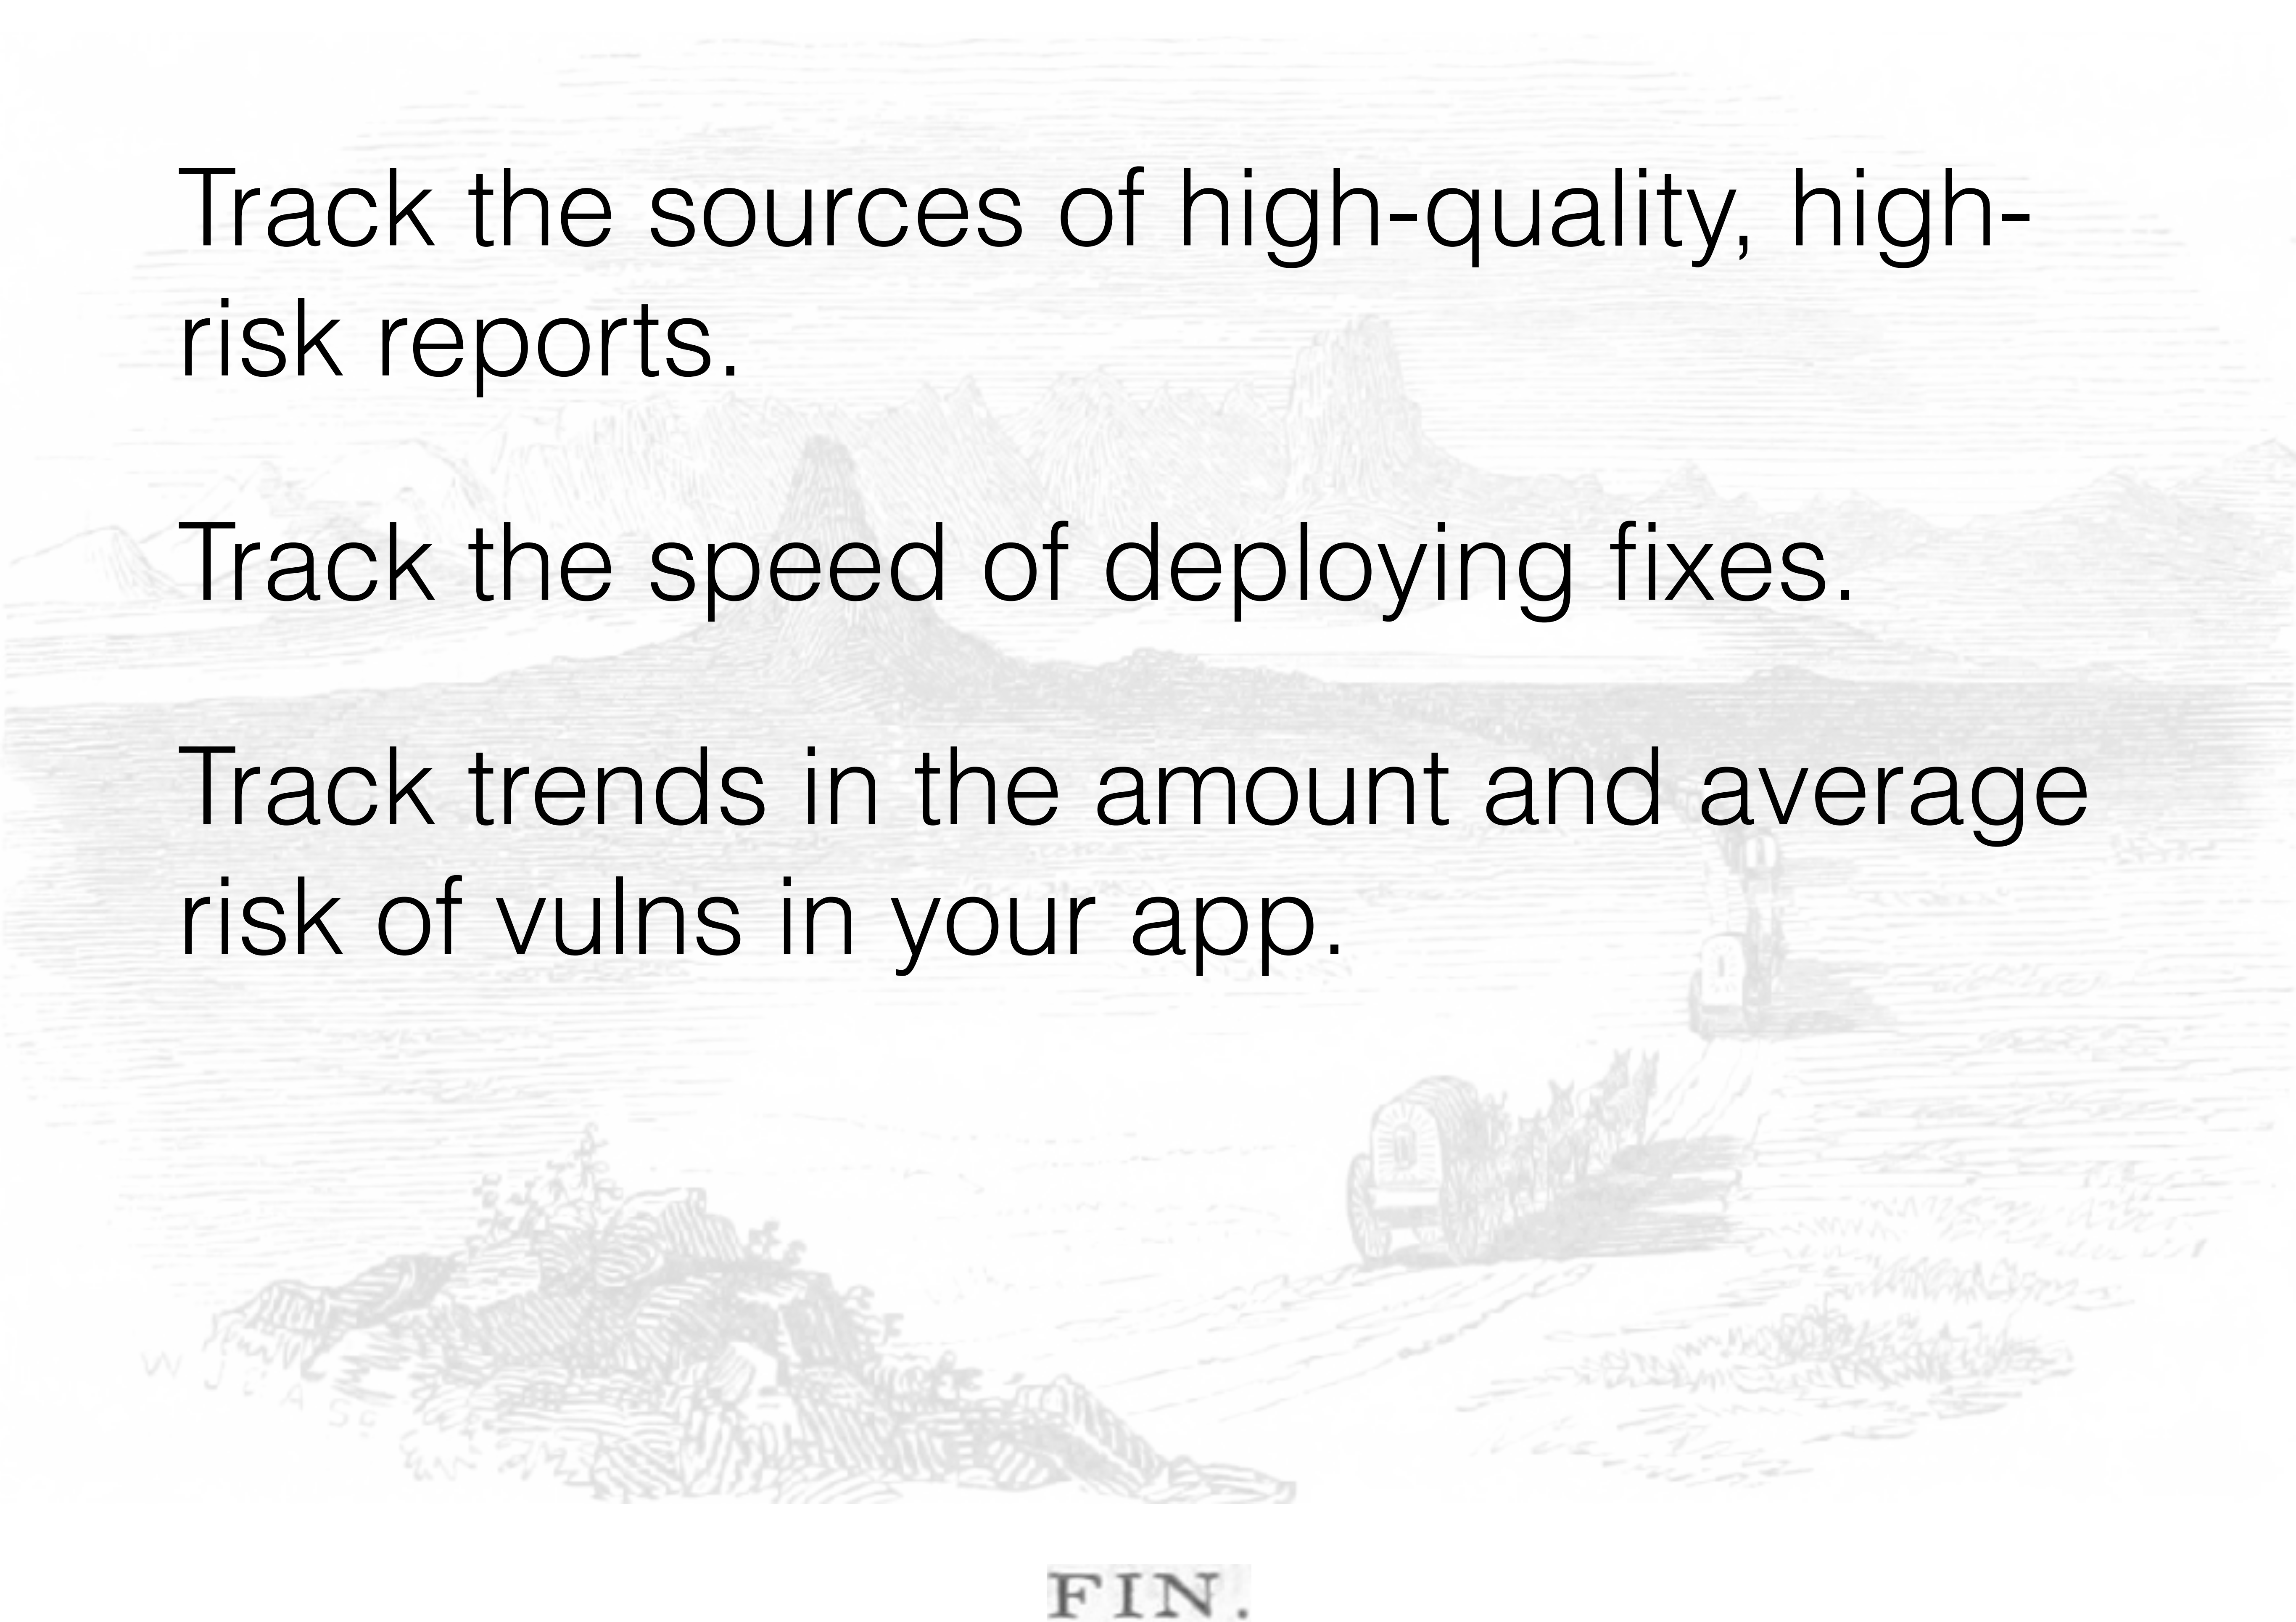

Thank You!

https://blog.cobalt.io

Questions?

In December 2013 the British Library released public domain images for anyone to use, remix, and repurpose.

Have fun!

https://www.flickr.com/people/britishlibrary/

Metrics: Program Value

Rate of incoming reports.

Time to mark valid/invalid (accept/reject).

Percentage of valid reports.

Price per valid reports.

Time to resolve valid reports.

Metrics: SDL Maturity

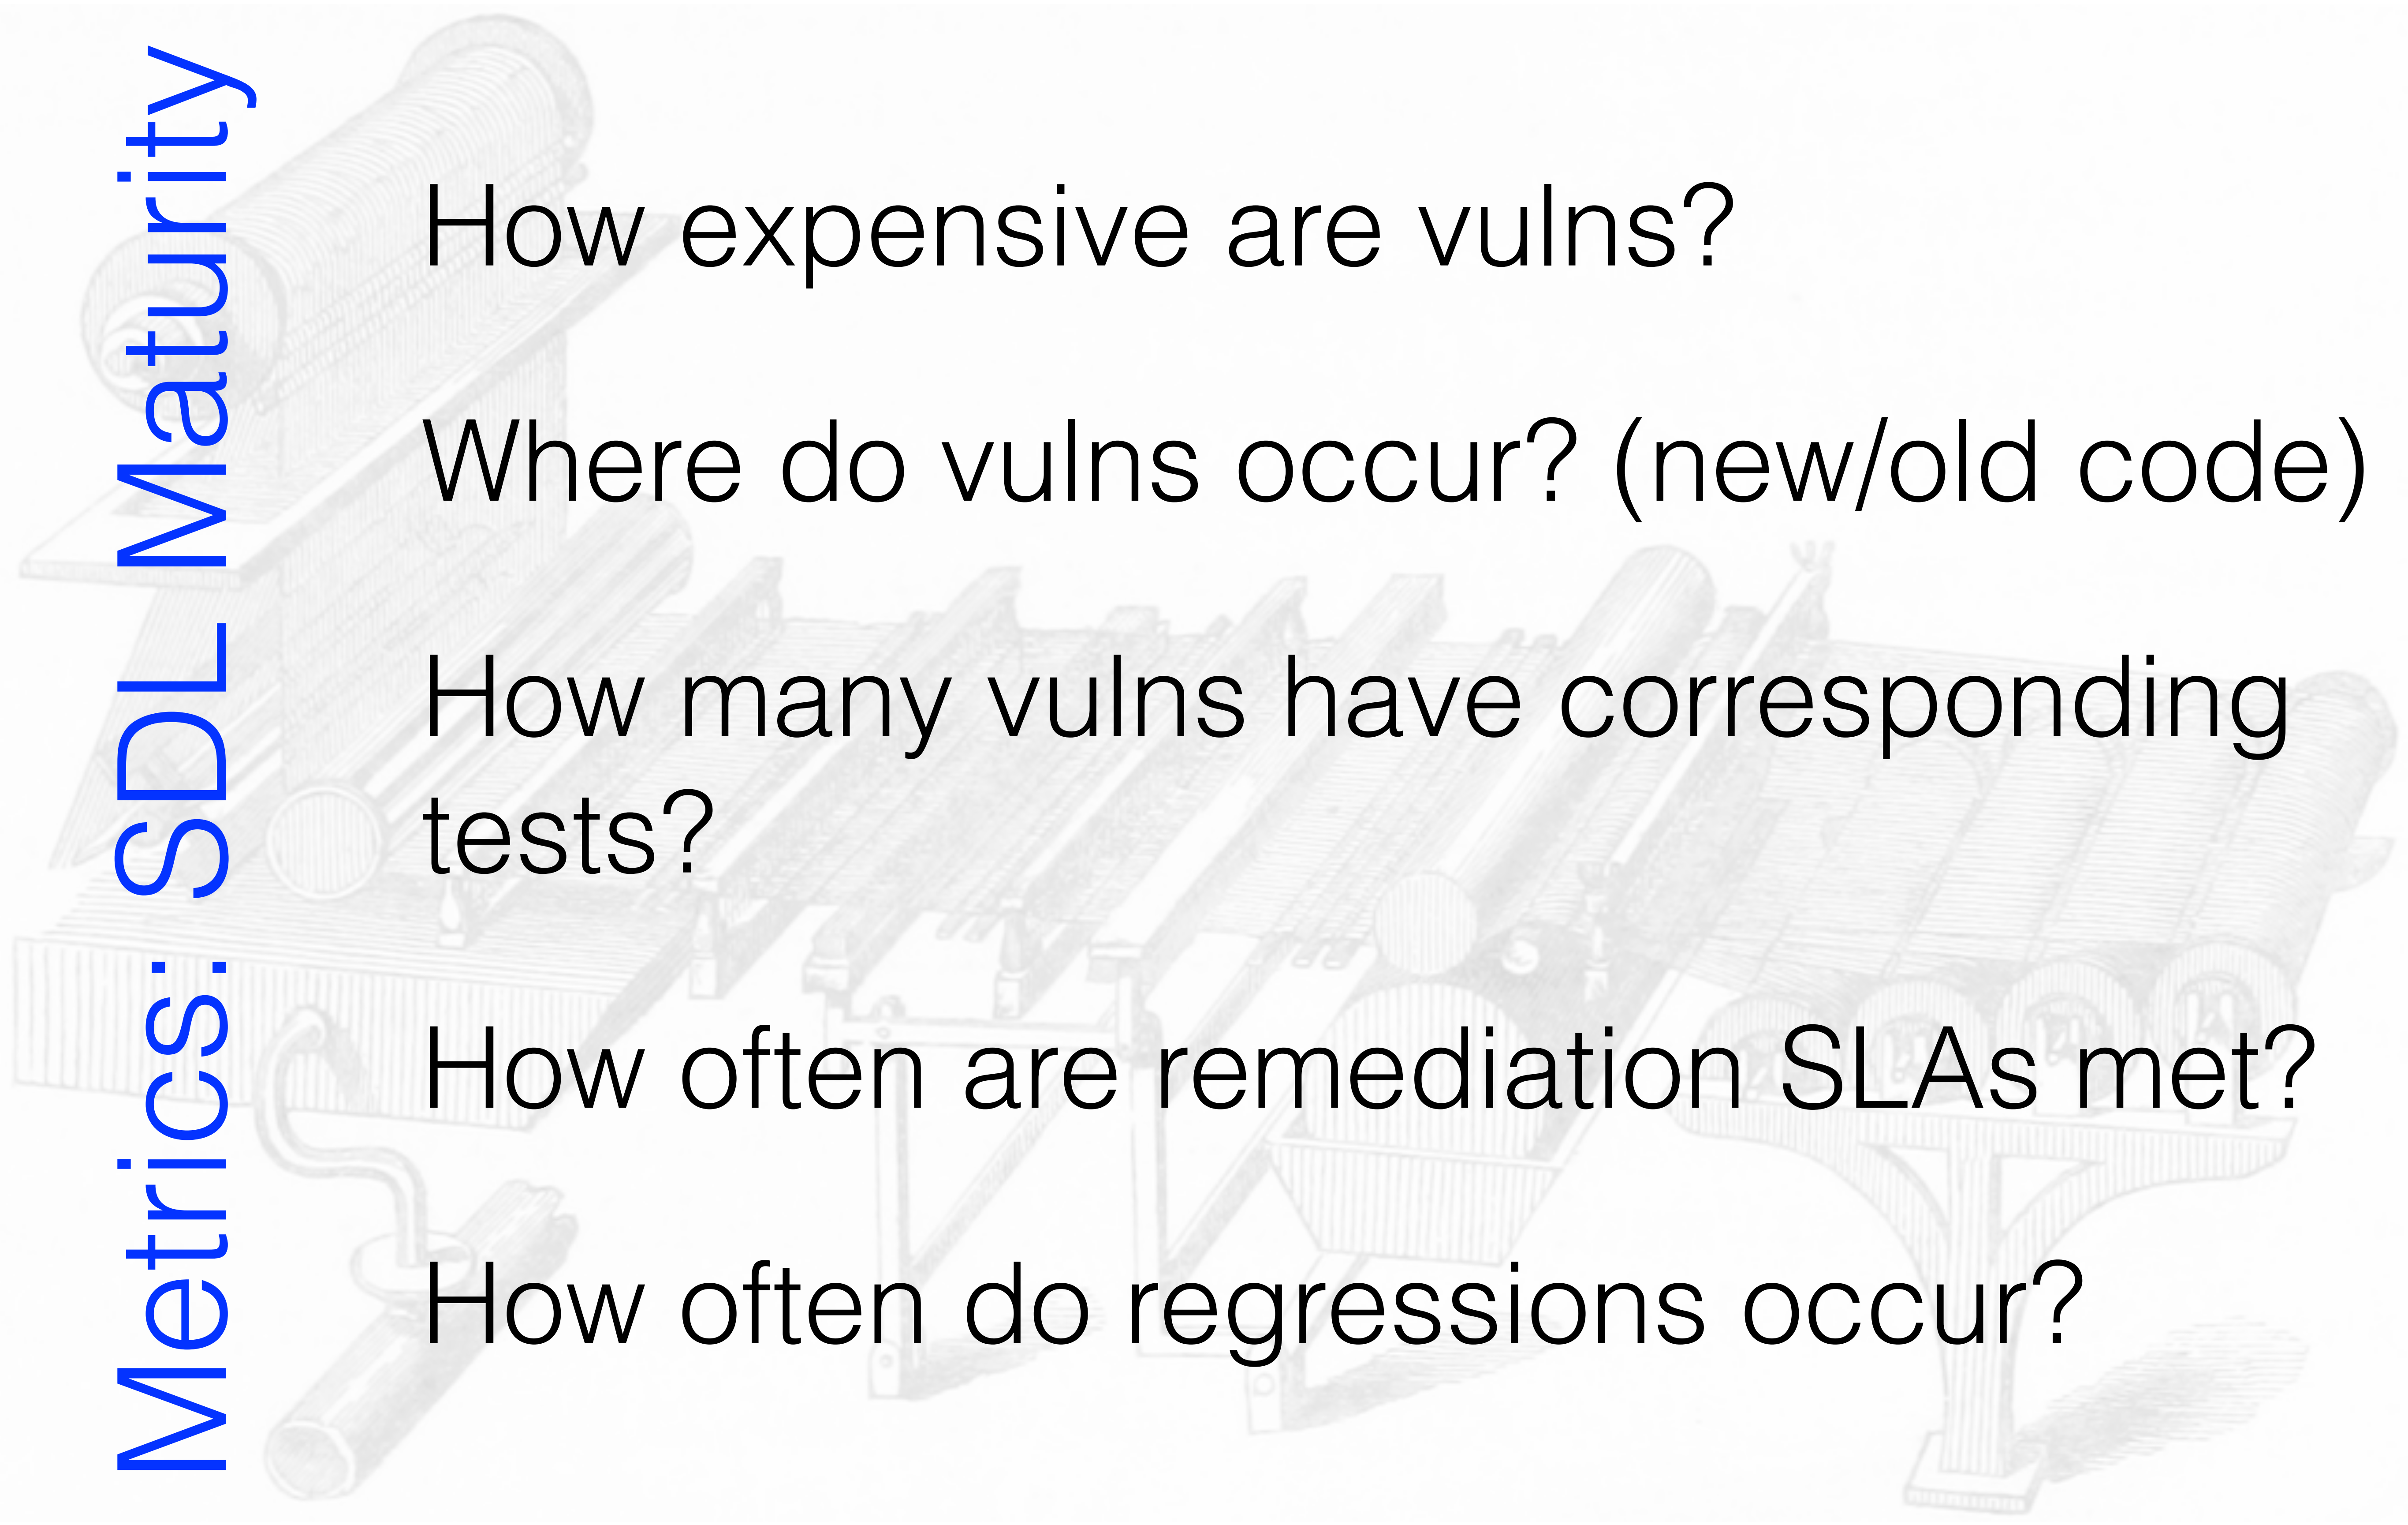

How expensive are vulns?

Where do vulns occur? (new/old code)

How many vulns have corresponding tests?

How often are remediation SLAs met?

How often do regressions occur?