From Chaos to Opportunities via security.txt

Matthew Dekker

OWASP New Zealand Day 2026

From Chaos to Opportunities via security.txt

ASMR Version

Matthew Dekker

OWASP New Zealand Day 2026

From Chaos to Opportunities via security.txt

Matthew Dekker

OWASP New Zealand Day 2026

Matthew Dekker

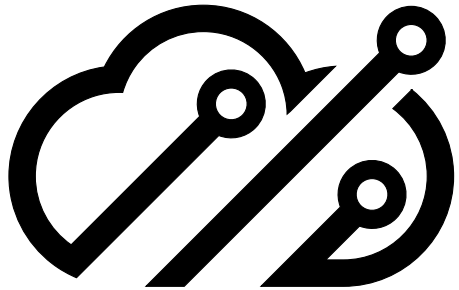

PRIVSECCONSULTING

- Principal Consultant & Practice lead for technical testing at PrivSec Consulting
- Focused on Application Security
- Almost lost my foot around this photo ->

Disclaimer

Some chaos was involved in constructing this talk so….

- The views in this talk don't represent my employer.
- All security research was performed in my own personal time out of view of my employer.

Intro

security.txt

But Vulnerability Disclosure Policies (VDP)s in general and a couple of war stories

Intro

You've heard of it, you probably know what it's for. But statistically speaking, you probably don't have one.

Intro

And if you do have one, is your whole process setup to encourage researchers while controlling all those beg bounties or AI slop reports?

What is security.txt?

What is security.txt?

It's a file on your website

mozilla.org/.well-known/security.txt

```
Email: security@mozilla.org
Main info: https://www.mozilla.org/en-US/security/
Bounty program: https://www.mozilla.org/en-US/security/bug-bounty/
```

google.com/.well-known/security.txt

```
Contact: https://g.co/vulnz
Contact: mailto:security@google.com
Encryption: https://services.google.com/corporate/publickey.txt
Acknowledgments: https://bughunters.google.com/
Policy: https://g.co/vrp
Hiring: https://g.co/SecurityPrivacyEngJobs
Expires: 2030-04-01T00:00:00z
```

apple.com/.well-known/security.txt

```
Contact: https://security.apple.com

# Apple Security Updates
Acknowledgments: https://support.apple.com/HT201222

# Apple Web Server Security Acknowledgements
Acknowledgments: https://support.apple.com/HT201536

# Apple Security Bounty Guidelines
Policy: https://security.apple.com/bounty/guidelines/

Expires: 2027-06-24T01:00:00.000Z
```

What is security.txt?

/.well-known/security.txt

or (if you have to drop in root)

/security.txt

What is security.txt?

It tells people how to contact you for security related things

Also where your VDP is and how you'll handle reports

What is security.txt?

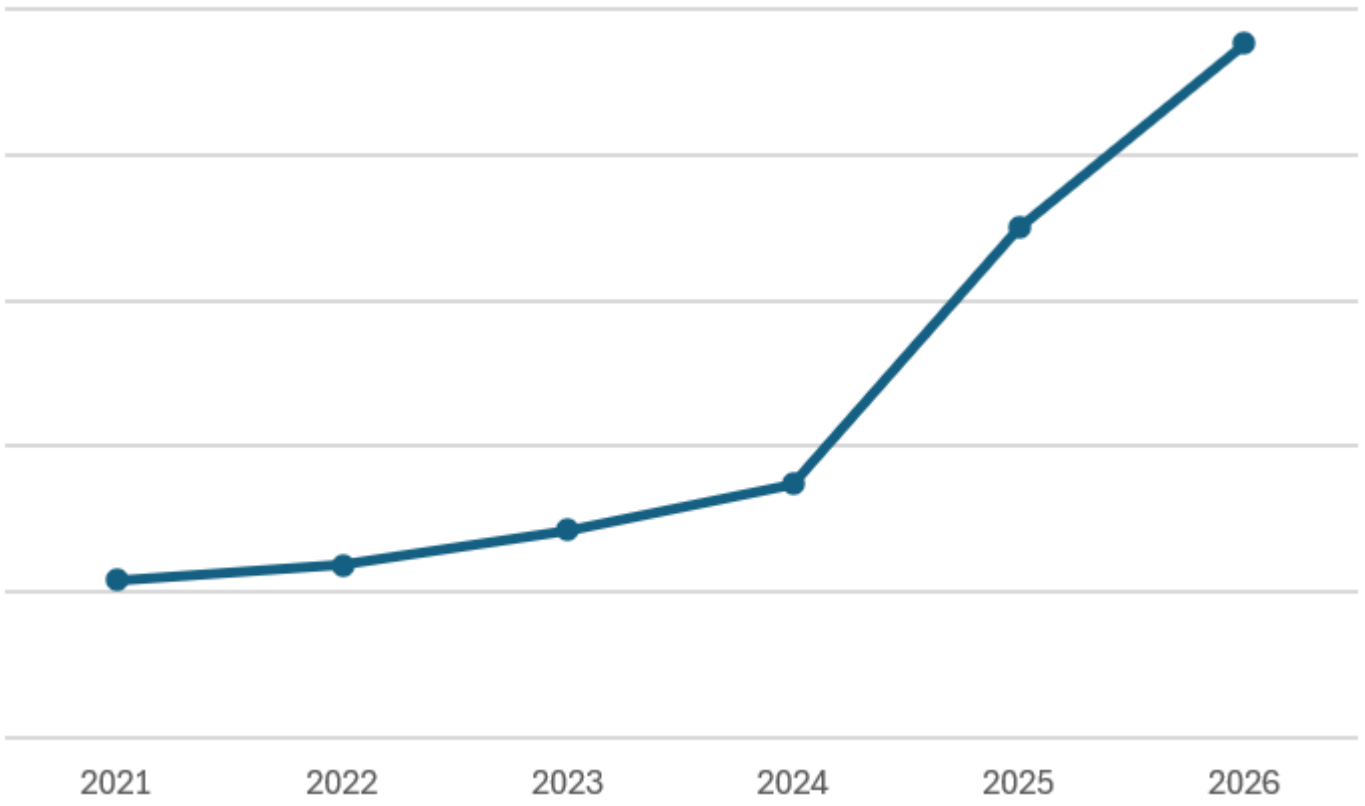

https://blog.iotdef.com/the-state-of-security-txt-adoption-an-analysis-of-240-million-domains-in-2026/

What is security.txt?

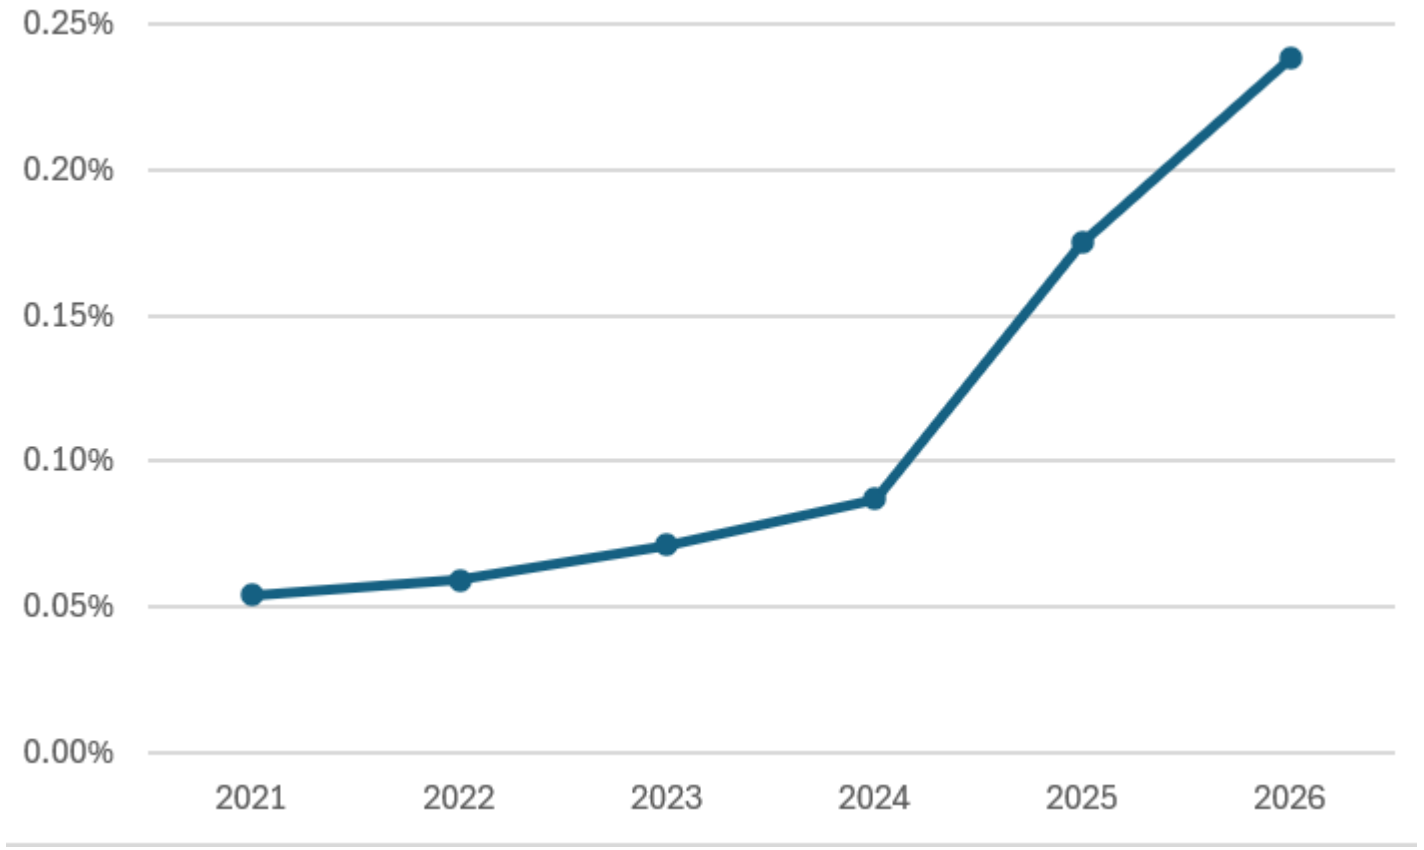

The Numbers

YEAR	DOMAINS SCANNED	WITH SECURITY.TXT	ADOPTION RATE
2021	151,402,656	82,343	0.054%
2022	171,101,802	100,950	0.059%
2023	193,282,908	137,320	0.071%
2024	210,286,332	183,767	0.087%
2025	223,234,116	390,085	0.175%
2026	241,285,150	573,123	0.238%

https://blog.iotdef.com/the-state-of-security-txt-adoption-an-analysis-of-240-million-domains-in-2026/

What is security.txt?

2022

15% of the Mozilla top 500

5% of the FTSE 100

3.6% of the S&P 500

2023

9.8% of the top 10k domains

2025

1.25% of the top 1 million domains

https://hexiosec.com/blog/survey-of-security-txt/
https://www.uriports.com/blog/security-txt-in-2025/

What is security.txt?

Most important:

Contact	Email address / Link to form

Secondary:

Policy	How you'll handle it & the legal stuff
Expiry	When the security.txt should be considered stale
Preferred-Languages	What language(s) you want to receive information in

Then think about:

Encryption	Where your PGP key lives
Acknowledgements	Hall of Fame for researchers who have reported or should otherwise be celebrated
CSAF	Common Security Advisory Framework or how you'll triage report severities

What is security.txt?

Other methods:

- Github projects
- Vulnerability Disclosure Programs
- DNS Security TXT

In all cases, without this info:

- People don't know where to go
- You can't control the conversation and expectations for security reports

What is security.txt?

OWASP has a Vulnerability Disclosure Cheat Sheet which goes into various aspects of the process.

Vulnerability Disclosure Cheat Sheet

Introduction

This cheat sheet is intended to provide guidance on the vulnerability disclosure process for both security researchers and organizations. This is an area where collaboration is extremely important, but that can often result in conflict between the two parties.

What is security.txt?

Security.txt is mentioned in regards to both finding and publishing contact details:

Finding Contact Details ¶

The first step in reporting a vulnerability is finding the appropriate person to report it to. Although some organizations have clearly published disclosure policies, many do not, so it can be difficult to find the correct place to report the issue.

Where there is no clear disclosure policy, the following areas may provide contact details

- Bug bounty programs such as BugCrowd, HackerOne, huntr.dev, Open Bug Bounty o Standoff.
- A security.txt file on the website at `/.well-known/security.txt` as per RFC 9116. At min include the required Contact and Expires fields; optional fields such as Preferred-Langua Canonical, Policy, and Hiring improve routing for ethical researchers.

Publishing Contact Details ¶

The most important step in the process is providing a way for security researchers to contact your organization. The easier it is for them to do so, the more likely it is that you'll receive secur reports. The following list includes some of the common mechanisms that are used for this - th more of these that you can implement the better:

- A dedicated security contact on the "Contact Us" page.
- Dedicated instructions for reporting security issues on a bug tracker.
- The common `security@` email address.
- A security.txt file on the website at `/.well-known/security.txt` as per RFC 9116. At mini include the required Contact and Expires fields; optional fields such as Preferred-Langu Canonical, Policy, and Hiring improve routing for ethical researchers.

Controlling the Chaos

Controlling the Chaos

Contact – Email address / Link to form

Controlling the Chaos

It's just a form/email address. But:

- Directing people where to go goes both ways:
  - You can tell them where they report
  - That method has to work

- It means you keep reports where they can be triaged and dealt with – keeping issues away from normal support processes.

- You're also promising that it'll be dealt with properly

Controlling the Chaos

You solemnly promise that:

- The contact method will work.
- The report will get in front of someone who can triage it.
- Responses will be prompt (you can define this expectation as well!)
- You will follow your own policy.

Types of Actors

SECURITY

Angry bug hunter with Microsoft beef drops new Windows 0-day

Revenge is a dish best served code

Jessica Lyons

62

Published Wed 10 Jun 2026 // 18:45 UTC

They are angry at Redmond and will have their revenge. Nightmare Eclipse, the prolific bug hunter and possibly disgruntled ex-Microsoft employee, disclosed another zero-day vulnerability just hours after Redmond issued a record-breaking number of CVEs and fixes for June Patch Tuesday.

Types of Actors

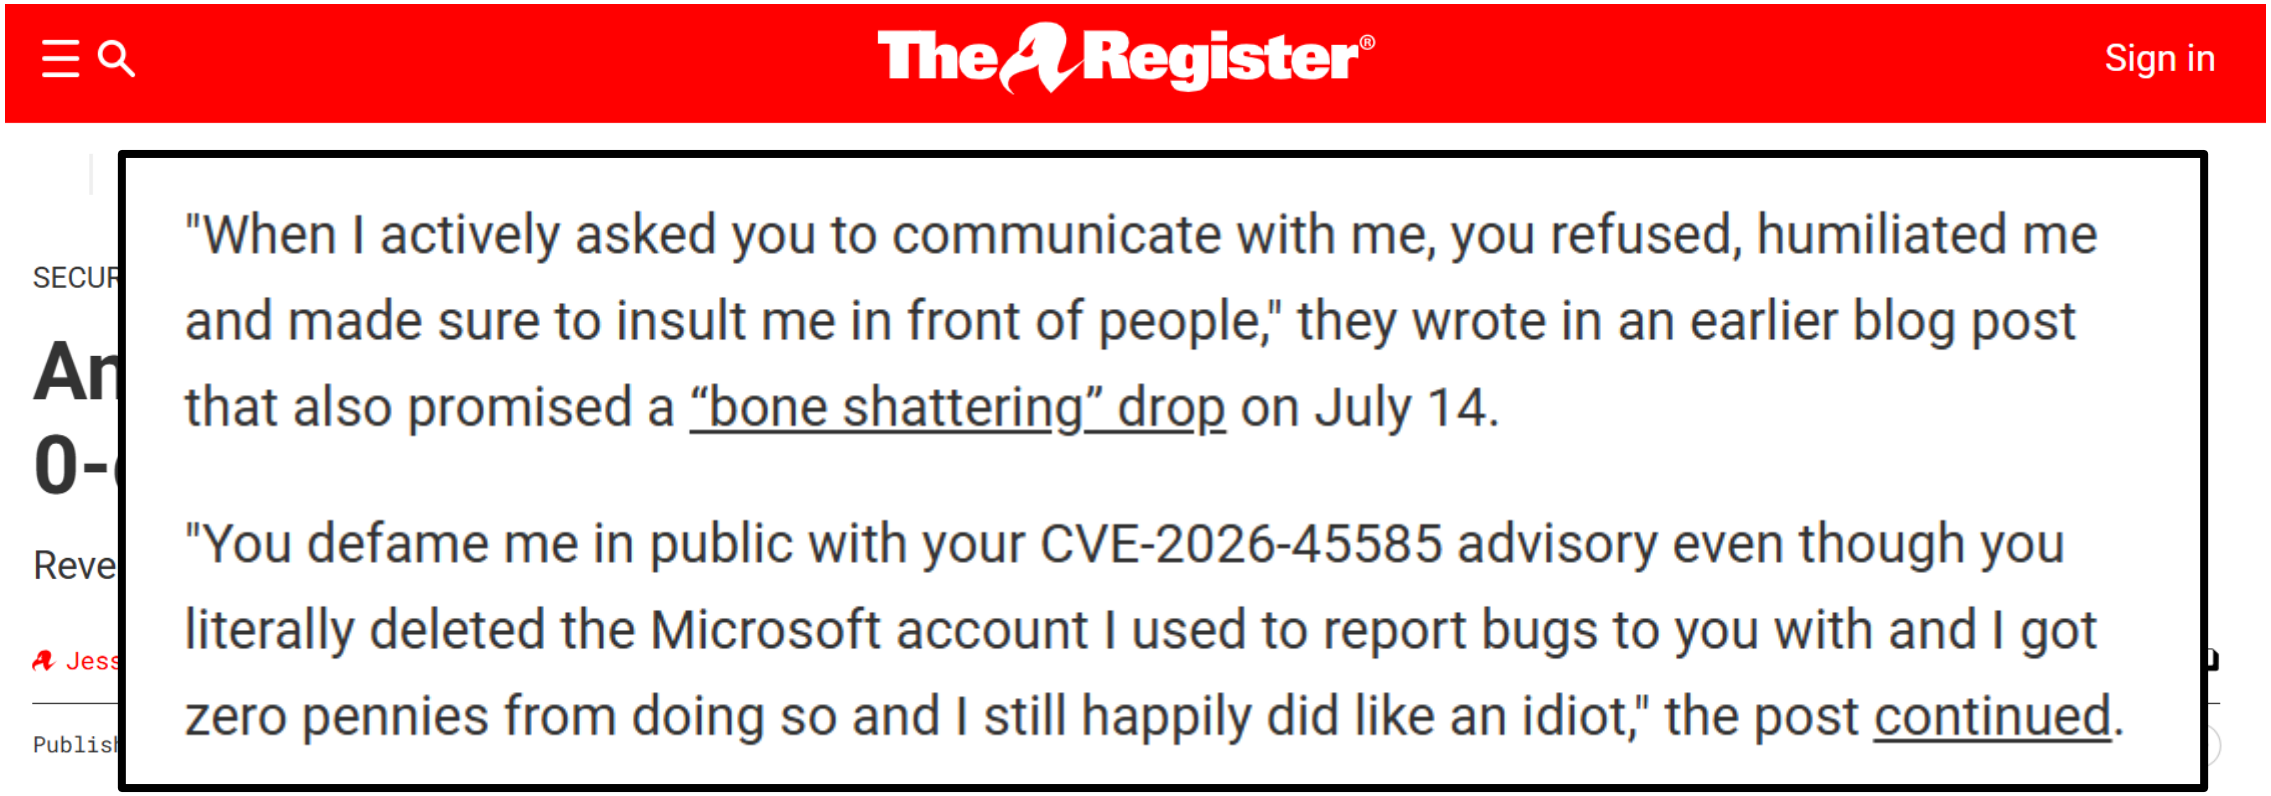

Controlling the Chaos

Quick tangent – even the big players get this wrong

Controlling the Chaos

Bug Bounty Platform triager didn't understand the app enough to triage the submission properly

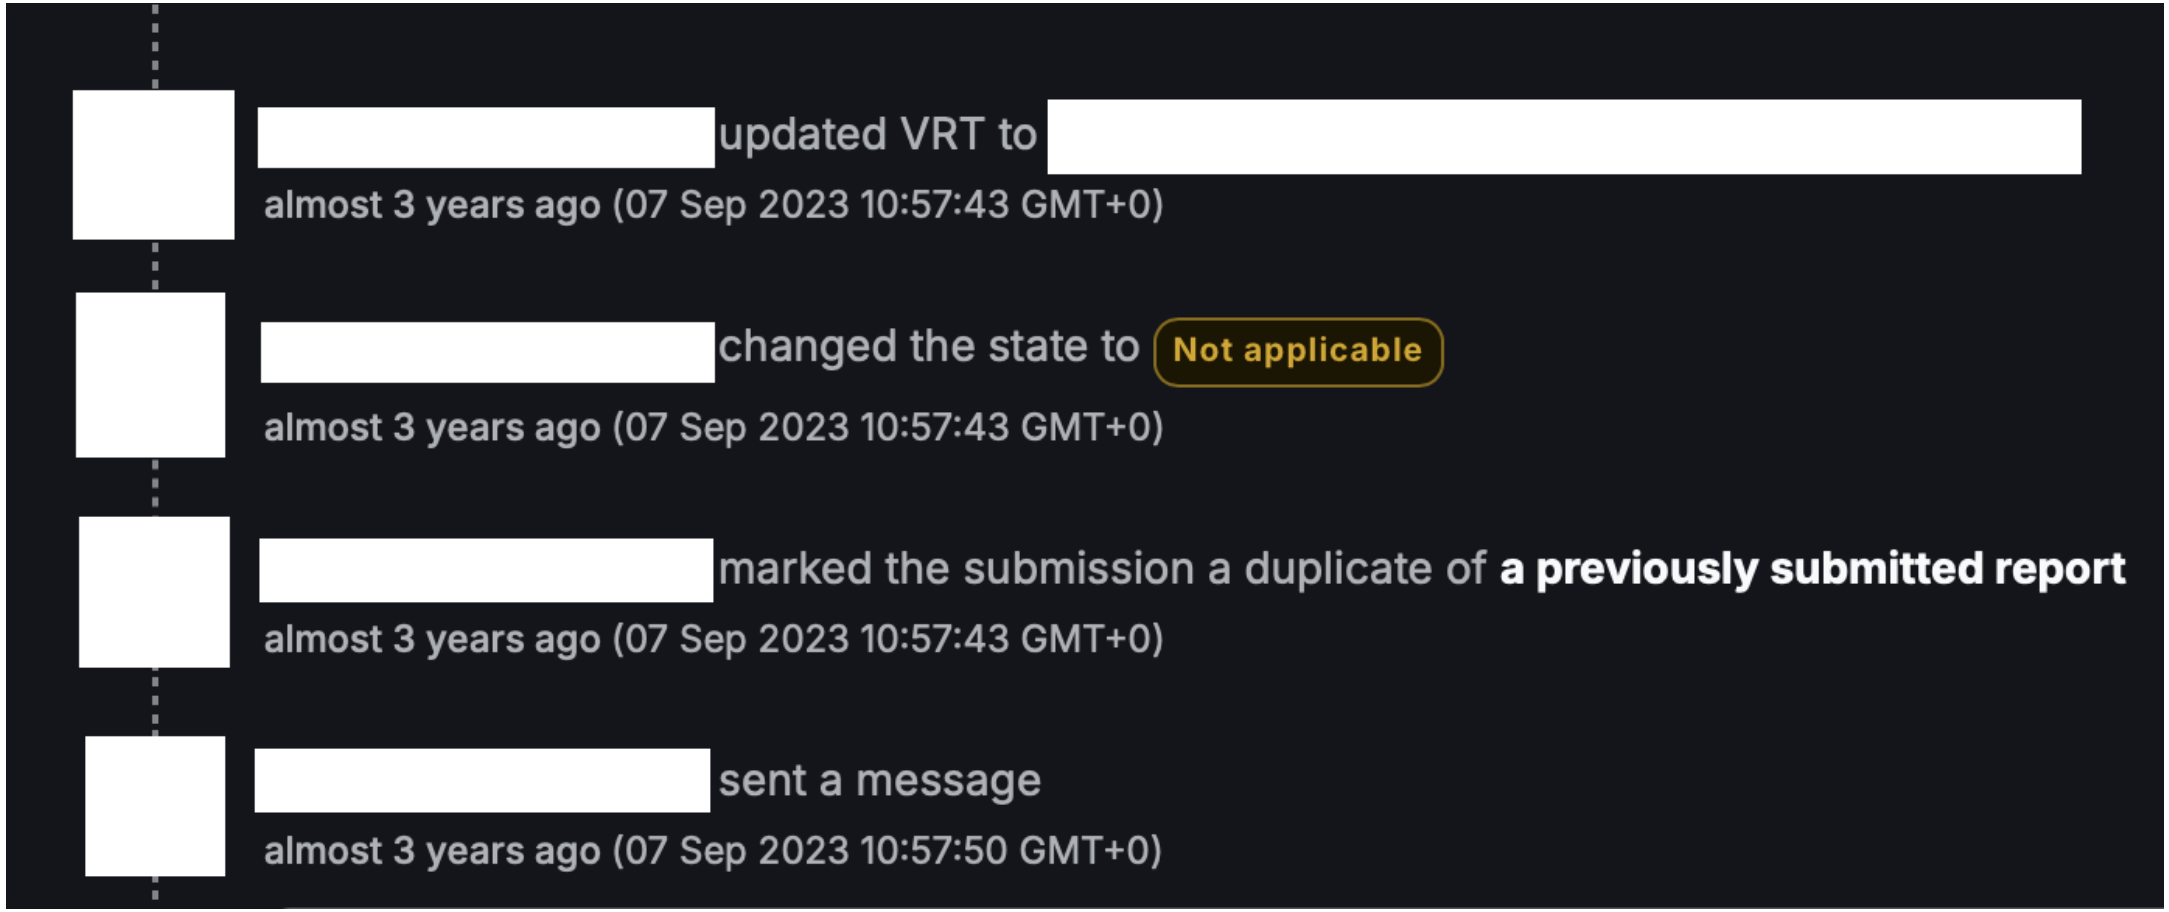

Controlling the Chaos

Resulting in an out of band email that needed to be dealt with

Matt Dekker Sent on: Fri, 08/09/23, 0:17

Concern over triaging

To

Hi there,

bit concerned that a recent bug might not have been triaged properly. It was marked as Duplicate and Not Applicable despite it being

Now, duplicate reports happen. My concern comes from the following points:

Sorry for reaching out to an old email, in my experience it's nearly impossible for a researcher to get a new response but I don't want to see this missed if it's not on your queue to fix.

Controlling the Chaos

Causing someone to have to manually step in

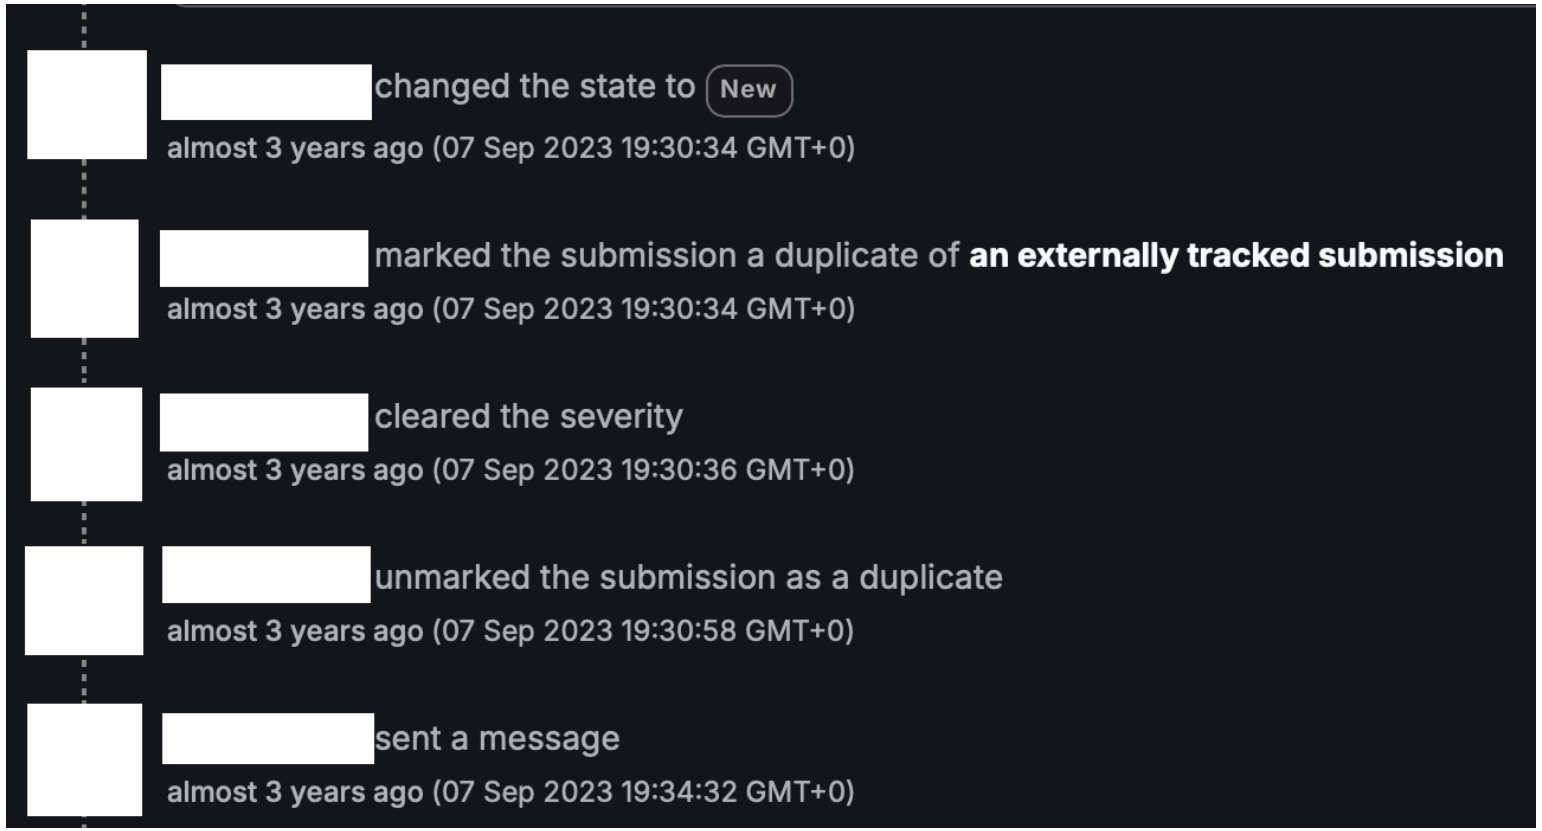

Controlling the Chaos

Before it got dealt with properly

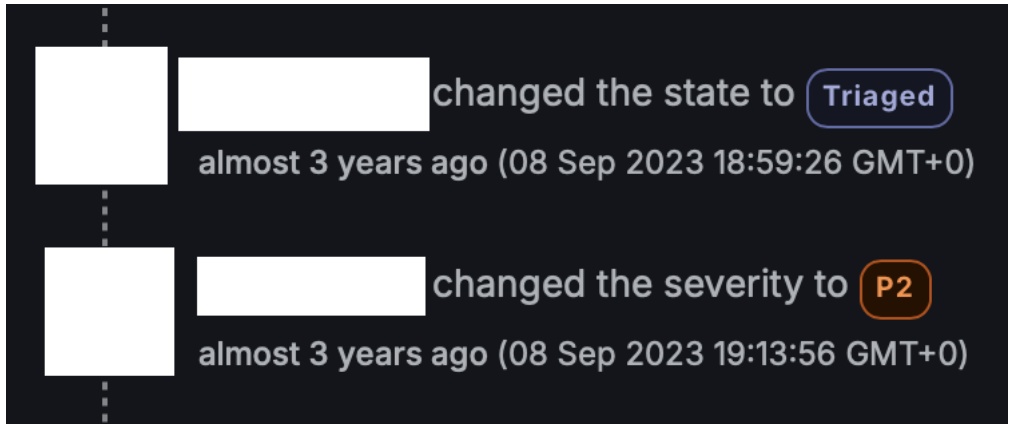

Controlling the Chaos

In summary, triage is hard.

Researchers can end up knowing more about your app than everyone except maybe your senior devs.

The key metric for fast resolution is getting reports through to someone who can deal with them.

Controlling the Chaos

Policy – How you'll handle it

Controlling the Chaos

This is essentially your public policy for how people can try to hack you. There can be legal implications, i.e. you accidentally authorise someone to DDoS your site as a "test".

Can include:

- Rules of Engagement
- Scope
- Vulnerabilities you care about…. And those that you don't
- Timelines & Responsible Disclosure Policies
- Safe Harbour
- Rewards, or the lack of them

Other sections as well, depending on what you want to cover.

Controlling the Chaos

Rules of Engagement

How you want people to behave when performing security research. If they break these, they won't be covered by any safe harbour. Can include:

- Not targeting other users without their explicit permission.
- Testing only against set environments and scopes.
- Adding a header to all requests.
- Not deleting any data.
- Reporting in a timely manner without abusing any vulns.
- No DoS / Phishing / Physical attacks / Illegal activities.

Controlling the Chaos

Scope

- The things you want tested
- The areas you explicitly don't want tested
- Areas of interest where you might provide a higher reward

Controlling the Chaos

Severity classifications & Vulnerability Exclusions

- Severity -> Reward classifications
- Usually a severity classification framework such as CVSS or the Bugcrowd VRT
- Things you don't want to hear about:
  - Security headers
  - DKIM & SPF misconfigurations
  - Theoretical SSL/TLS bugs without a PoC
  - Out of date JS libraries

Controlling the Chaos

Timelines & Responsible Disclosure Process

How quickly you'll aim to:

- Respond
- Triage
- Fix
- Pay any rewards

What you'd like researchers to follow in terms of public disclosure. Standard is 90 days from report before going public

- This is weird as technically researchers don't "need" to play nice and can talk about their work but you can threaten to not pay any rewards if they breach this.
- Can get awkward when the clout from a disclosure is viewed as worth more than your reward scheme.

Safe Harbour

Explicit permission to perform security research against targets – intended to avoid legal repercussions.

This is probably the biggest carrot you have to attract researchers.

- Legal provisions such as safe harbor policies.
  - The disclose.io project provides some example policies.
  - Take legal advice from lawyers, **not from this cheat sheet**.

Controlling the Chaos

Safe Harbour

https://disclose.io includes wording such as the following:

When conducting **vulnerability research according to this policy**, we consider this research conducted under this policy to be:

- **Authorized in view of any applicable anti-hacking laws**, and we will not initiate or support legal action against you for accidental, good faith violations of this policy;
- **Authorized in view of relevant anti-circumvention laws**, and we will not bring a claim against you for circumvention of technology controls;
- **Exempt from restrictions in our ToCs** that would interfere with conducting security research, and we waive those restrictions on a limited basis; and
- **Lawful, helpful to the overall security of the Internet, and conducted in good faith.**

You are expected, as always, to **comply with all applicable laws**. If legal action is initiated by a third party against you and you have complied with this policy, **we will take steps to make it known that your actions were conducted in compliance with this policy**.

If at any time you have concerns or are uncertain whether your security research is consistent with this policy, please submit a report through one of our Official Channels before going any further.

Controlling the Chaos

From the researcher's view – no safe harbour means significant risk to doing research.

Controlling the Chaos

I am a user of your platform who also does security research. I wanted to reach out and see if you had a policy in place regarding safe harbour for any security related issues and reports for the platform. If so, what would be the best way to report anything that were to come up?

To be clear, I don't have anything to report, just want to clarify if this is something that you'd be open to receiving.

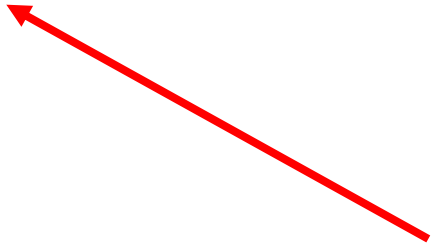

Probably not a lie. But also not entirely true.

Controlling the Chaos

Rewards, or the lack of them

- Whatever you're offering, state it here.
- If you're not offering anything, state that.

This is a set of expectations around what the researcher can expect in exchange for spending time on your app. Whatever you set here should be met or exceeded.

Controlling the Chaos

Minimal examples:

Security.txt:

Contact: security-vdp@example.com

Policy: https://example.com/security-policy.txt

This has legal implications and I am not a lawyer.

Security Policy ideas:

- Scope – this website
- Out of scope
  - Physical attacks, DoS, Social engineering, illegal actions
  - Security headers, SSL/TLS, Exploits without impact, Out of date software without a PoC
- Responsible Disclosure – You'll fix it or request an extension within 90 days and want to coordinate a disclosure with the researcher.
- Safe Harbour – There is permission to hack the website as long as the rules are followed.
- Rewards – Our eternal thanks, no bounties.

Many other things can be added but the above is a decent start

Controlling the Chaos

Free Research?

- Safe Harbour – There is permission to hack the website as long as the rules are followed.
- Rewards – Our eternal thanks, no bounties.

Controlling the Chaos

Remember these numbers?

The Numbers

YEAR	DOMAINS SCANNED	WITH `SECURITY.TXT`	ADOPTION RATE
2021	151,402,656	82,343	0.054%
2022	171,101,802	100,950	0.059%
2023	193,282,908	137,320	0.071%
2024	210,286,332	183,767	0.087%
2025	223,234,116	390,085	0.175%
2026	241,285,150	573,123	0.238%

Just having a security.txt already makes you more attractive to security researchers than 99.7% of websites

https://blog.iotdef.com/the-state-of-security-txt-adoption-an-analysis-of-240-million-domains-in-2026/

Types of Actors

But why are we inviting people to hack us?

Types of Actors

The ones who ignore rules will do it regardless. You may as well create a route for someone who finds something and wants to report it.

Types of Actors

Types of Actors

Types of Actors

Intent

Cause harm	Have Fun
Make money	Make Things Better

Types of Actors

Intent	General response to a process	Notes
Cause harm	Neutral	Won't care
Make money	Slightly positive	Wants the path of least resistance to be paid.
Have fun	Positive	More reason to keep within bounds, has an easy reporting channel to discourage public "naming and shaming".
Make things better	Very Positive	Unlikely to engage without a process. Might lead to ongoing research.

Types of Actors

- The good guys will know where to go and feel like you'll take their report seriously.
- Those looking to spam will hopefully go to the right place rather than whatever they can find.
- The bad guys weren't looking for your permission anyway.
- And in all cases, you can point to something to structure the conversation.

Types of Actors

Those looking to spam will hopefully go to the right place rather than whatever they can find.

AI being the big concern here these days

Hey ChatGPT,

Find all security issues in this website and send them a bug bounty report with maximum impact asking for a bounty.

Types of Actors

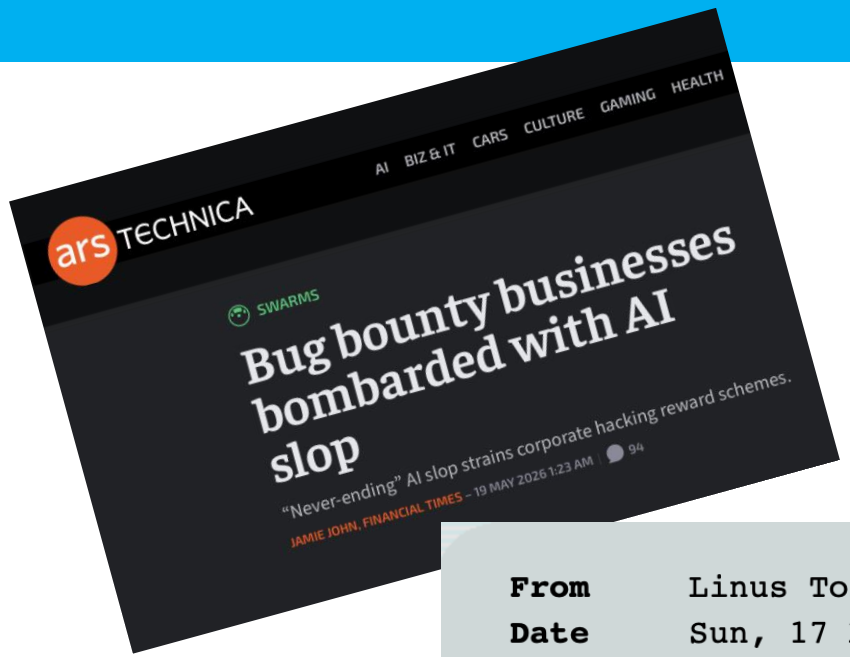

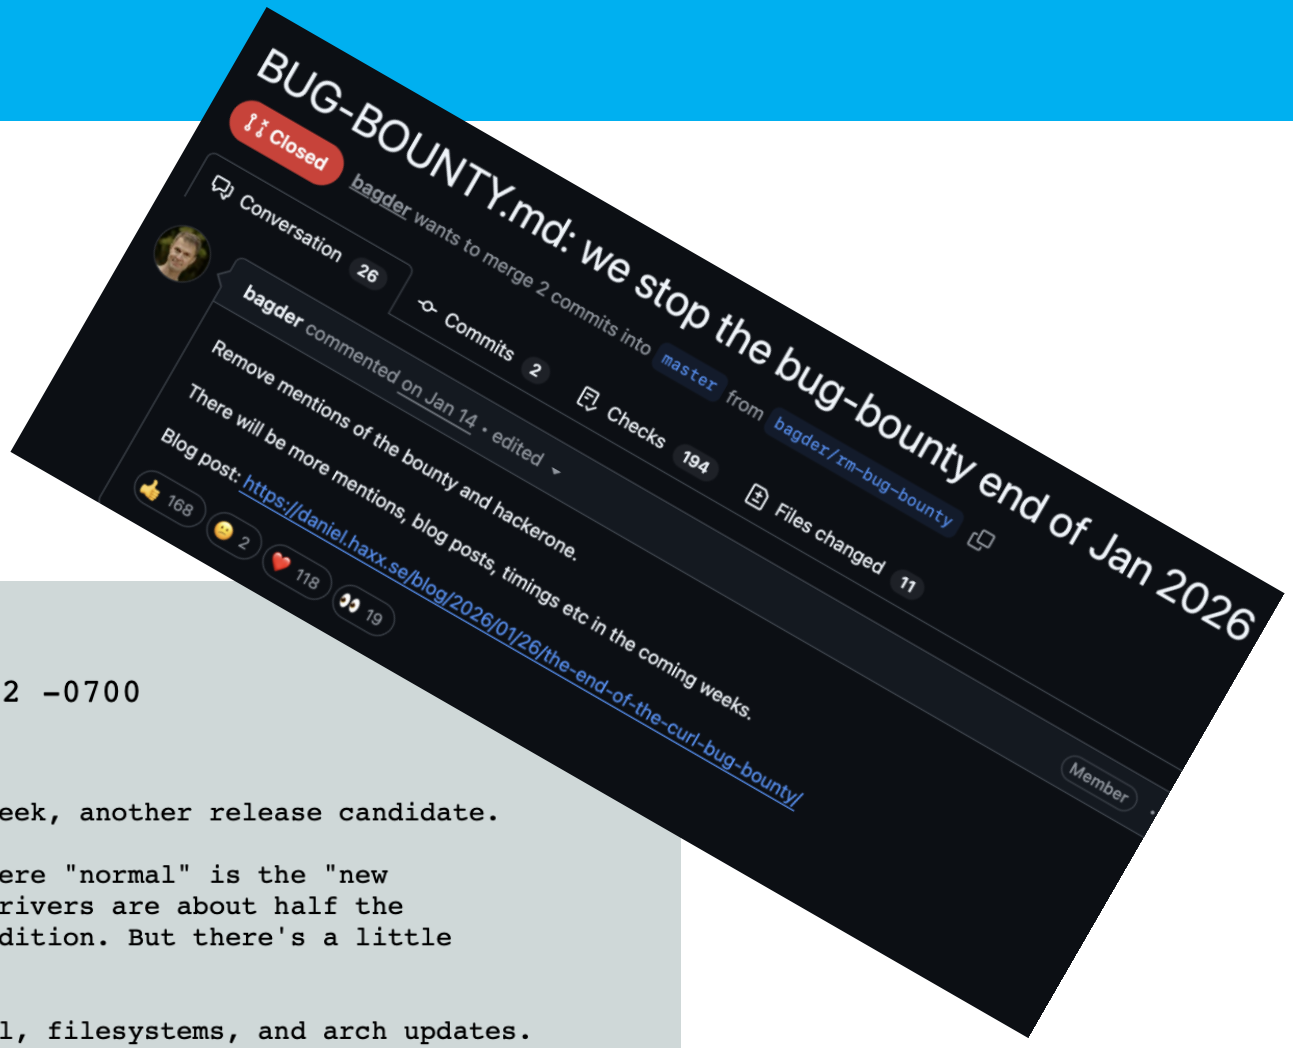

From Linus Torvalds <>
Date Sun, 17 May 2026 14:29:22 -0700
Subject Linux 7.1-rc4

You all know the drill by now - another week, another release candidate.

Things continue to look fairly normal (where "normal" is the "new normal" with a fair amount of changes). Drivers are about half the patch, with GPU leading the way as is tradition. But there's a little bit of everything in driver land.

The rest is mostly networking, core kernel, filesystems, and arch updates.

Some of the documentation updates might be worth highlighting: the continued flood of AI reports has basically made the security list almost entirely unmanageable, with enormous duplication due to different people finding the same things with the same tools. People spend all their time just forwarding things to the right people or saying "that was already fixed a week/month ago" and pointing to the public discussion.

Types of Actors

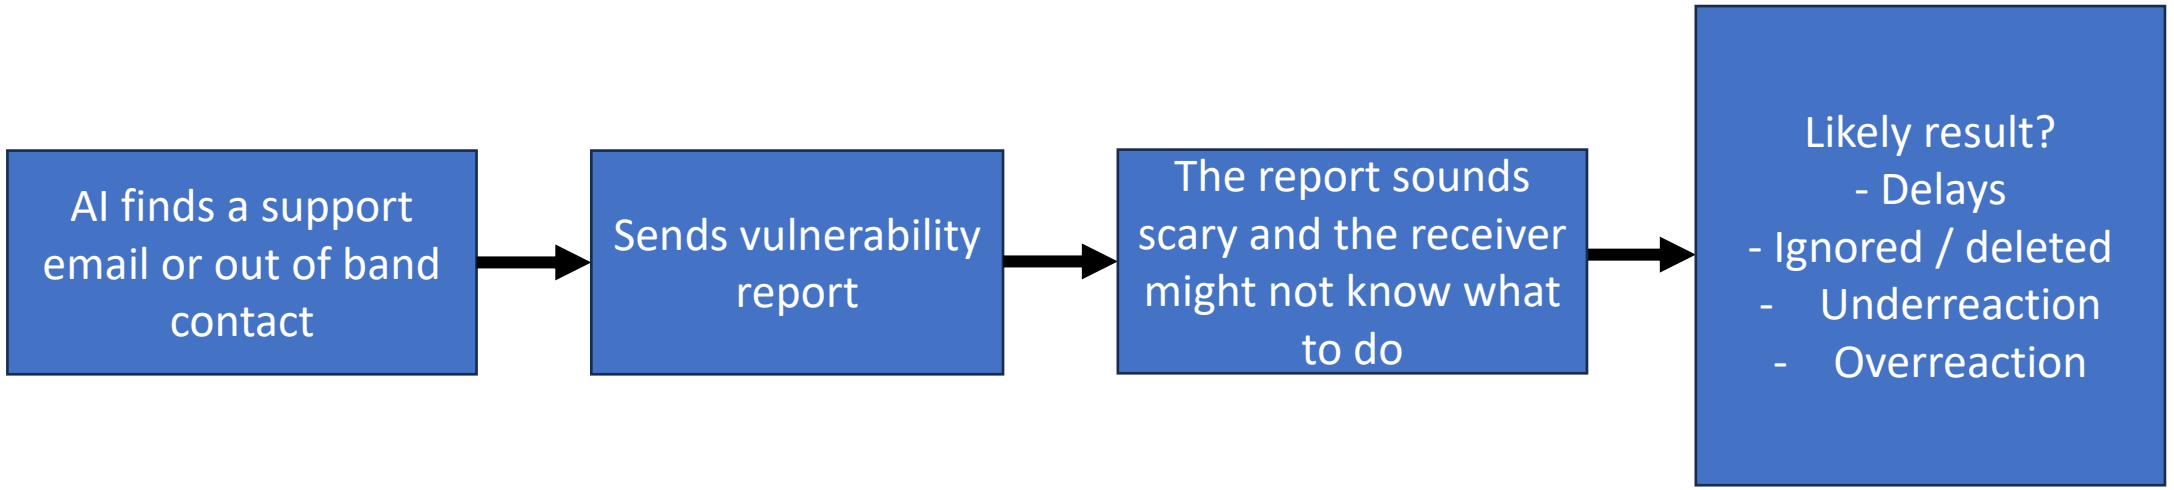

Types of Actors

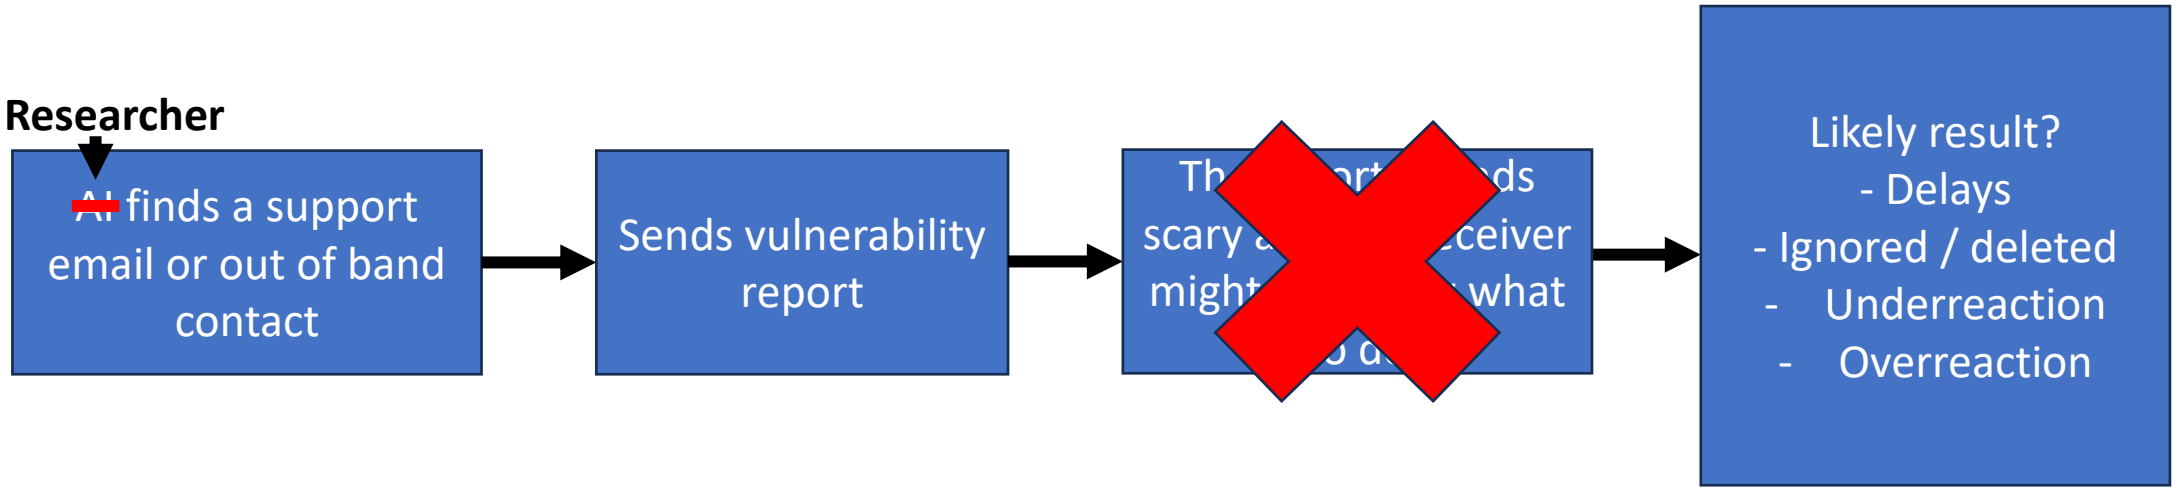

And researchers love to complain if they feel like they weren't listened to.

Types of Actors

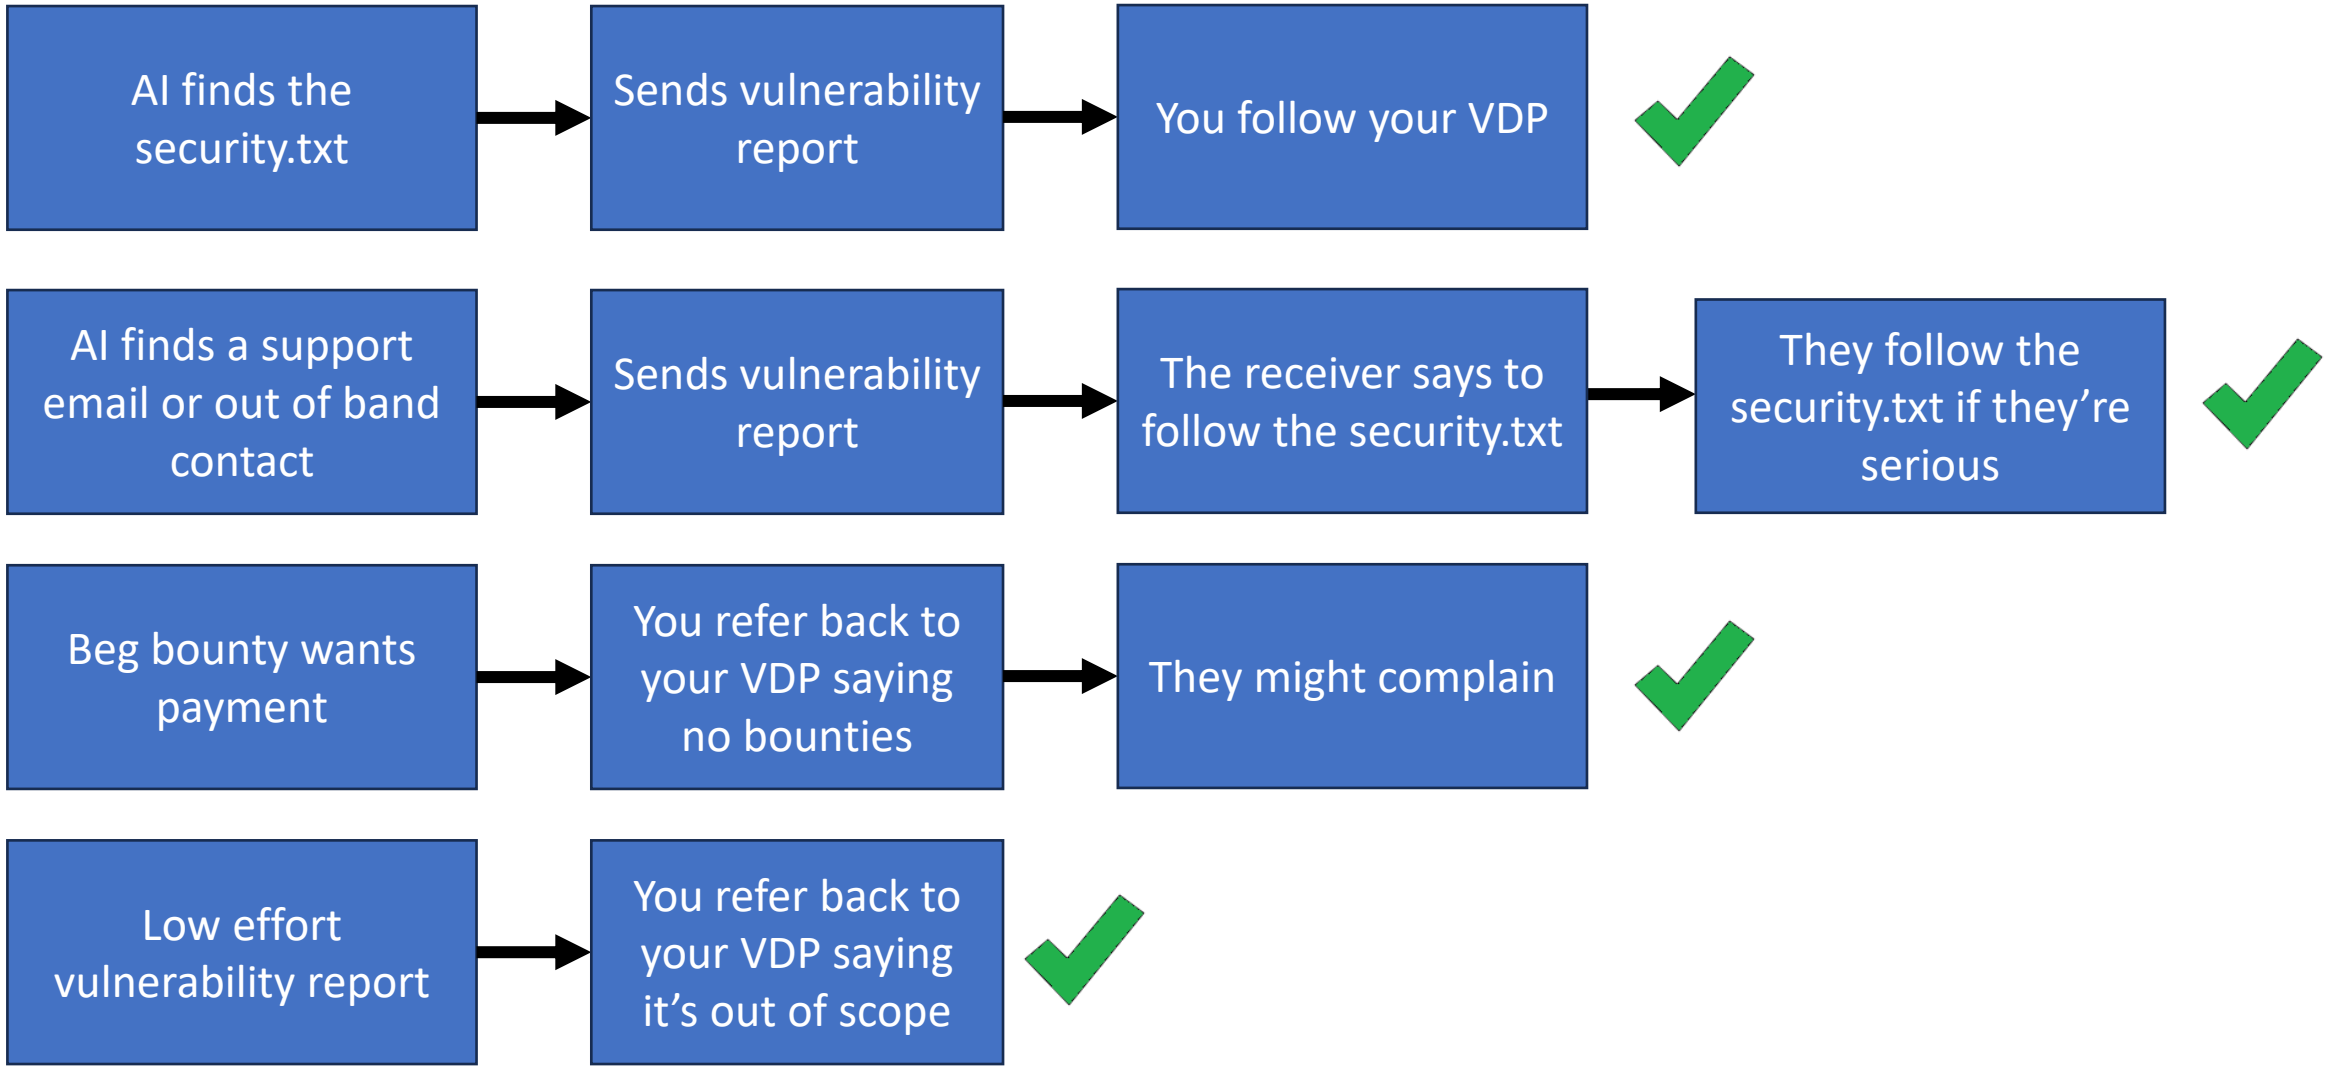

Types of Actors

- The people behind your VDP become a point of failure.
- Since they can triage the reports, they're probably quite busy already.
- Adding your own pre-parse with AI or a more junior staff member to filter the garbage and ensure that all required information has been received can go a long way.
- Can you ensure that the reporting flow can still be completed in a timely manner during holidays?

Beg Bounties

Example time!

Beg Bounties

While at work, a client had a beg bountier in their emails. They sent through an anonymised transcript for this talk:

Hi Company,

I'm Beg Bounty Hunter from BegBountyHunters Inc. I'm reaching out because we recently identified a critical security vulnerability on the Company platform.

Given your 20-year history architecting software and data systems—from your early work on the Previous Work platform to leading Company today—I wanted to flag a critical security vulnerability I came across today.

Beg Bounties

The vulnerability we identified could allow an attacker to gain access to users' credentials, personal data, and cause major reputational damage. We

We're happy to give you more details about the vulnerability before we proceed any further. BegBountyHunters Inc. will keep your file for Company for a 15-day period, so please respond to us at your earliest convenience.

Beg Bounties

- **Fee Structure:** Our fee for this report is $2,550 USD. We only send an invoice once the vulnerability has been fixed.

For ONE vulnerability!

Probably was just a dangling DNS record in AWS which required winning the IP lottery.

Bounties

Singular vulnerabilities do fetch large sums..... If the company operates a program to pay them.

- In this case, the likely vulnerability was quickly identified and resolved.
- The bug class wouldn't have been easily exploitable.
- A VDP indicating that the company wouldn't pay would have likely stopped the "report" happening, saving time and stress.

Per the OWASP Cheat Sheet:

- Do not demand payment or other rewards as a condition of providing information on security vulnerabilities, or in exchange for not publishing the details or reporting them to industry regulators, as this may constitute extortion.

Bounties

Bounties

- Bounties define what type of researcher you're going to get.
- Making these clear and then keeping to them prevents a lot of frustration for everyone.
- You can always reward higher when you receive something special.
- Having a defined & smooth reporting process can bring researchers above what they'd be interested in from a purely financial perspective.
- Breaking your own VDP expectations can undermine any bounties you offer.
  - There are a number of well paying bounty programs out there that people famously avoid.
- Fun rewards can be highly sought after.

Bounties

Money	Mindful of sanction lists, currently conversions, international transfers
Hall of fame	Make it public and easy to find
Swag	Postage costs and worldwide delivery, build up hype around it
Your products/services	Could be essentially free for you
Allowing public disclosure	Let the researcher brag a little, but also show that you cared about the issue.
Other	Get creative to say thanks! Things I've seen: • The researcher get flown out for a meet & greet. • Invites to special events at conferences. • Ongoing commercial relationships.

Example – The Forge

Example – The Forge

Quick note:

Everyone in the following example did the best they could. Was it perfect? No. That's why this talk exists.

But things were taken seriously and fixed for a good result in the end. Be nice.

Maybe even consider checking their product out if you play Tabletop RPGs – it's pretty cool.

Example – The Forge

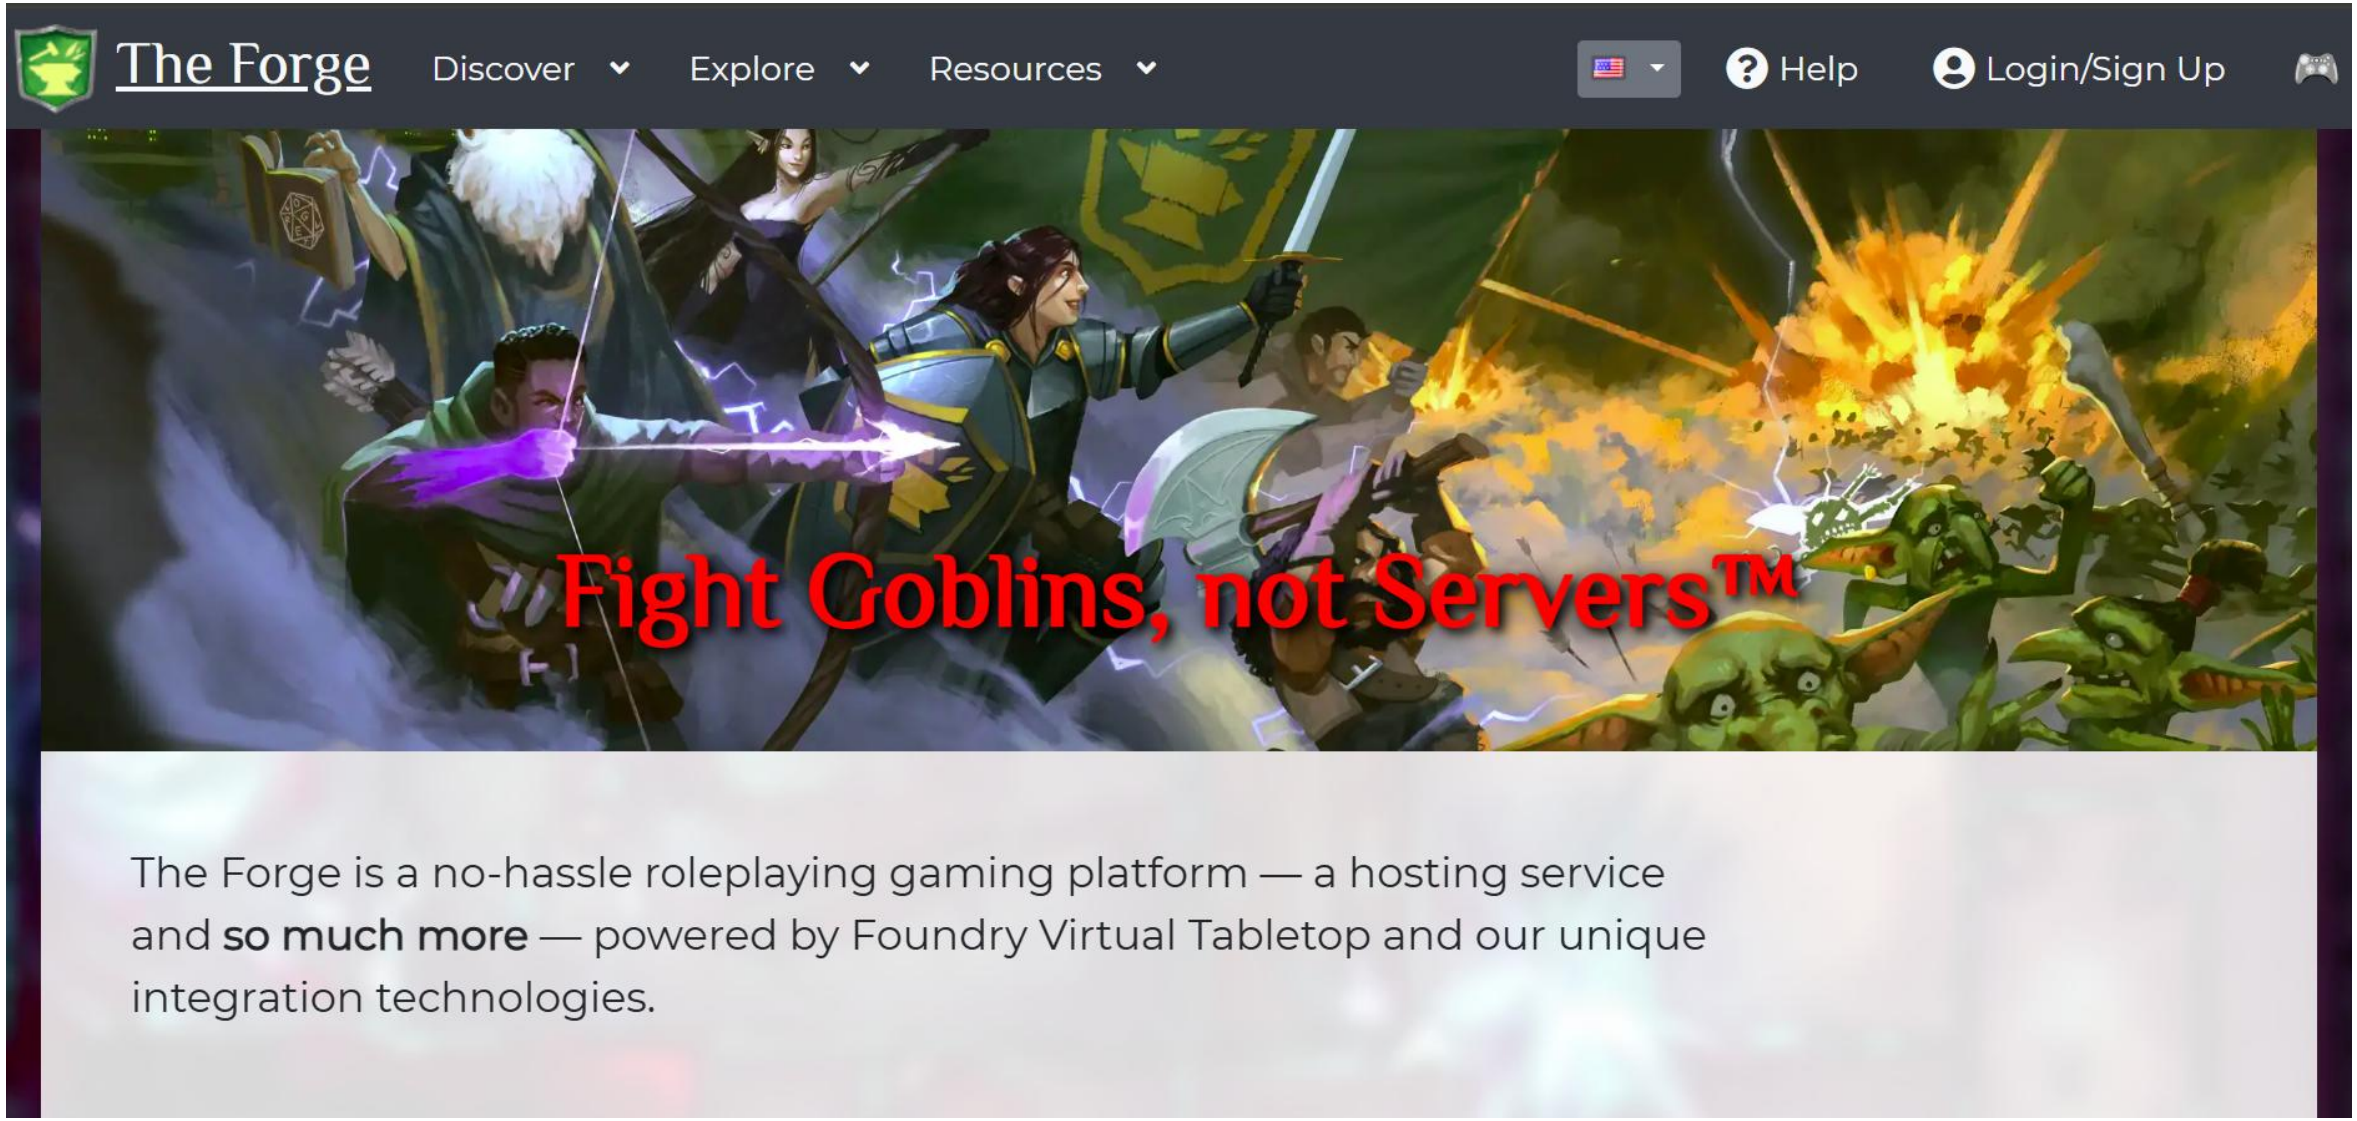

Example – The Forge

And their CEO was great about this whole situation:

As for your request, it would be my pleasure to give you authorization to disclose the name of the company in your talk. It sounds like a really interesting talk and would love to watch a recording of it if one of the local meetups records these types of events.
Also, I wanted to note how awesome it is that you're doing this and then sharing your experience and knowledge like that with your community, kudos to you!

Good luck for your talk, and don't hesitate to reach out if there's anything else I can help you with!

Thanks again,

Example – The Forge

Late Nov:

Signed up for a game through their platform.

- Thought it was interesting so looked for a security.txt or anything else to indicate safe harbour.
- Nothing.

Example – The Forge

December 2

Policy for Security Related Issues Inbox ×

Matthew Dekker < > Tue, Dec 2, 2025, 9:49 PM

to contact

Hey there,

I am a user of your platform who also does security research. I wanted to reach out and see if you had a policy in place regarding safe harbour for any security related issues and reports for the platform. If so, what would be the best way to report anything that were to come up?

To be clear, I don't have anything to report, just want to clarify if this is something that you'd be open to receiving.

Thanks for your time,
Matthew Dekker

Example – The Forge

December 13 – 11 days for response

Forge Support <contact@forge-vtt.com> Sat, Dec 13, 2025, 2:04 PM

to me

Hello Matthew,

Thank you for reaching out to us. We haven't dealt with a situation like this before, so we don't have any policies clearly written. If you are comfortable proceeding without such delineated policies, we would be happy to offer safe harbor for any ethical reporting of issues found on our platform.

Thank you,

The Forge

•••

Example – The Forge

Same day (already had some vulns)

Matthew Dekker < > Sat, Dec 13, 2025, 4:34 PM

to Forge

Hi

Thanks for the response. So, I took a look after your email and think I've got an account takeover issue, at minimum there's at least an XSS issue. I am about to jump into a game but can finish my PoC and send through details after.

Are you happy for me to send details unencrypted here or would you prefer me to use another mechanism?

Cheers,
Matt

•••

Example – The Forge

December 23 – 10 days to get a contact

Forge Support <contact@forge-vtt.com> Tue, Dec 23, 2025, 4:19 PM

to me

Hello Matthew,

You can send the security issue via security@forge-vtt.com.

Let us know if you have any other questions or concerns. We would be happy to help.

Thank you,

The Forge

Example – The Forge

December 30 – Emails from the beach

Matthew Dekker < > Tue, Dec 30, 2025, 8:00 PM

to Forge

Hi ,

That should be through to you now, please let me know if the email to security doesn't come through (it has a couple of attachments).

Thanks for your help here.

Cheers,
Matt

•••

Example – The Forge

Mail Delivery Subsystem <mailer-daemon@googlemail.com> Tue, Dec 30, 2025, 7:59 PM
to me

Address not found

Your message wasn't delivered to **security@forge-vtt.com** because the address couldn't be found, or is unable to receive mail.

LEARN MORE

Example – The Forge

Immediately followed up via Email but also opened Discord support ticket due to the holidays

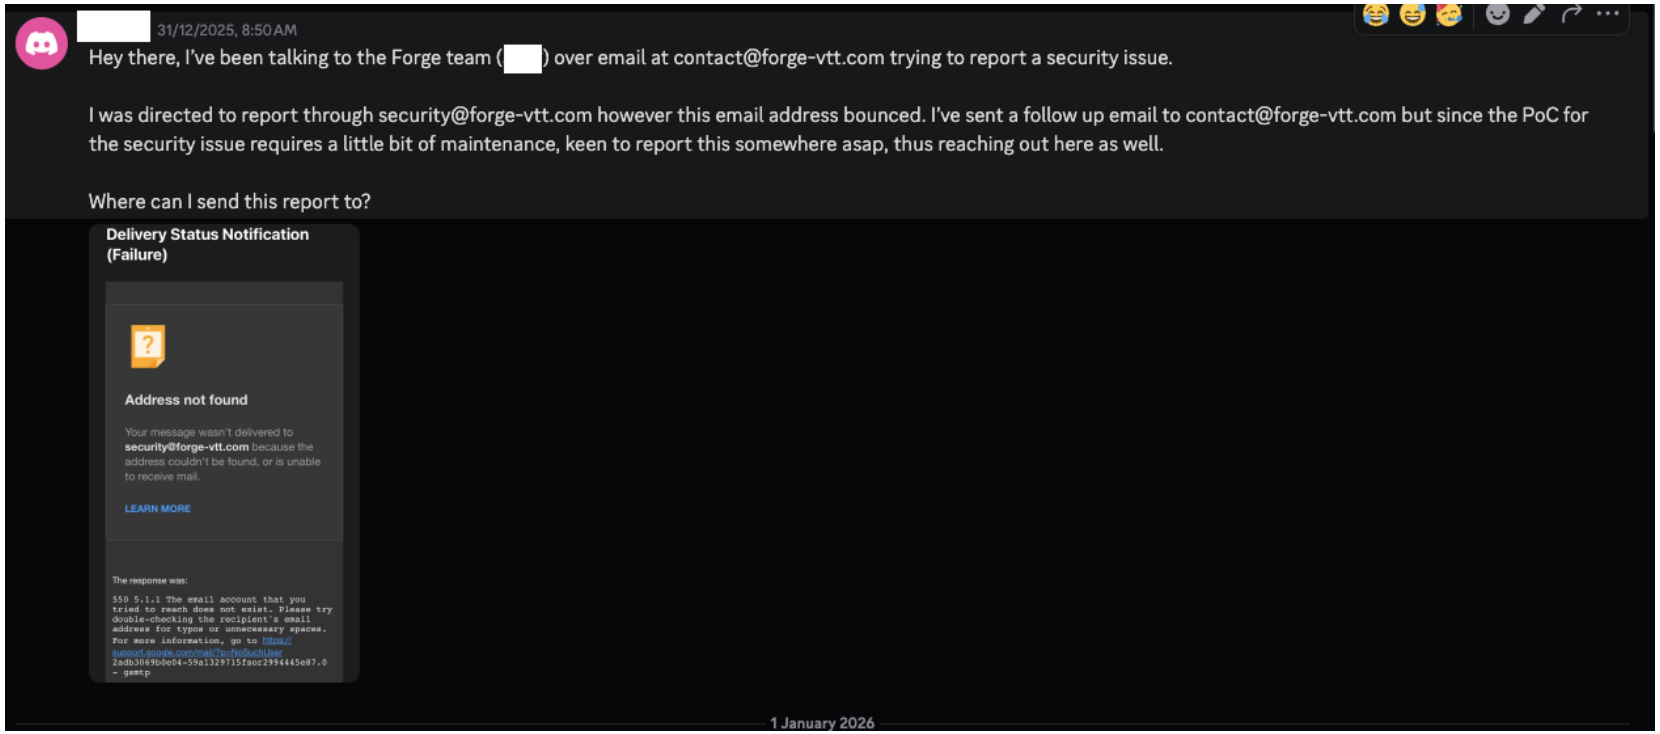

Example – The Forge

1 – 9 Jan, No progress

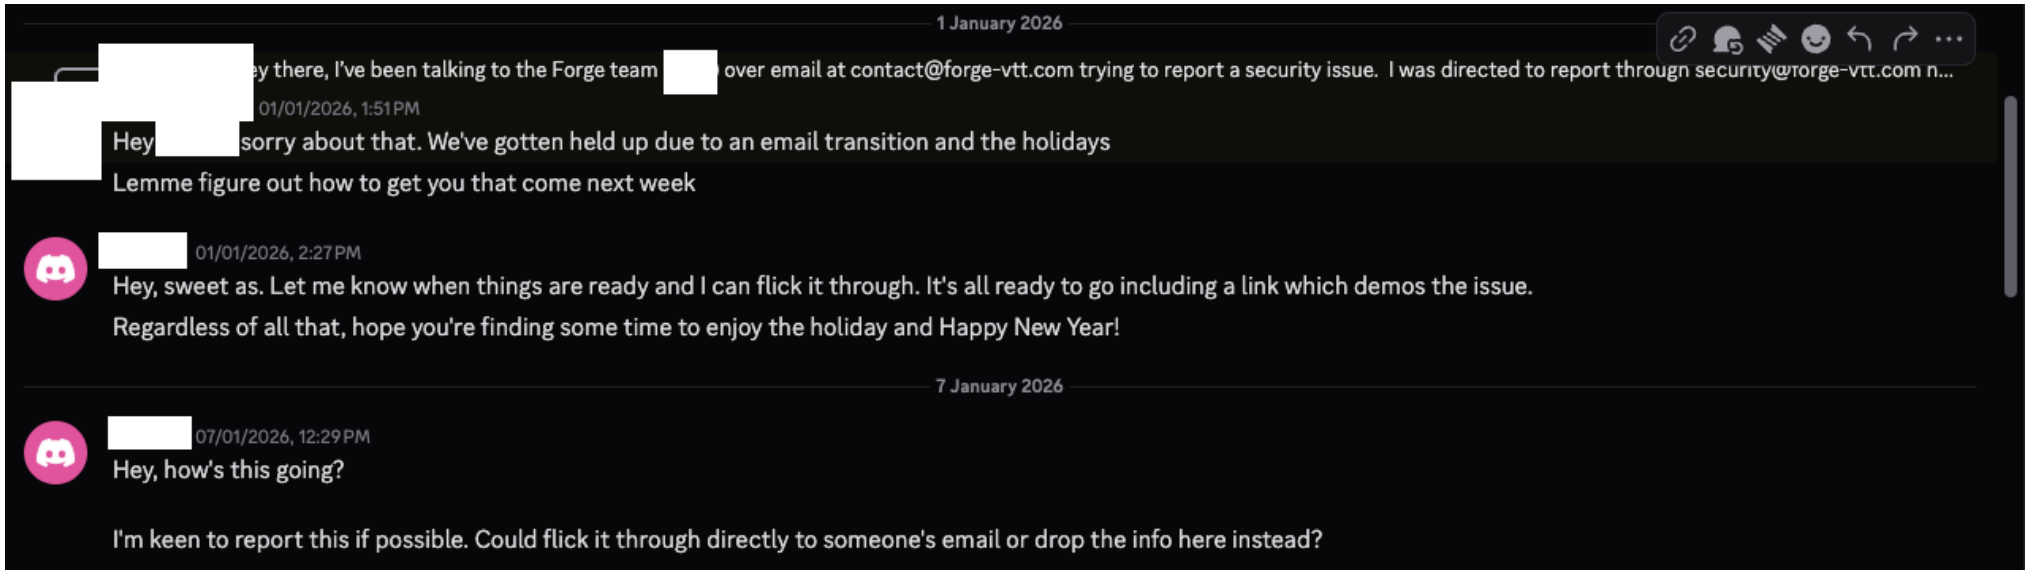

Example – The Forge

8 & 9 Jan, Anything for traction

Reporting a Security Vulnerability in Forge-VTT Inbox ×

Matthew Dekker < > Thu, Jar

Hi

Sorry to reach out to you completely out of band but I've been trying to report a security issue in
Holidays haven't really helped but still. I've been told to report to a security specific email but it v
bounce and a ticket (#2400) in Discord hasn't gotten me much further.

Not looking for a bounty or anything else, just want to report this thing and stop screaming into the void since it requires a little bit of backend maintenance. Any idea where I can send this report to?

Happy to drop it anywhere which is related to Forge-VTT (out of caution since this is an unverified channel).

Thanks for your time,
Matthew Dekker

Found a relevant staff member's personal email through OSINT

They were on leave but responded next day

[redacted] 09/01/2026, 10:34 AM

Hey,

Unless I hear otherwise, I'll be dropping the disclosure here in a few hours and also to any Forge-VTT emails that I've encountered.

Happy to follow a designated process but this has been right around a month now without being able to even report.

[redacted] since you're listed as [redacted] on the website, don't know if either of you would want jump in here directly?

Directly tagged devs then disclosed a few hours later

Matthew Dekker < > Fri, Jan 9, 1:14 PM

to **<Various Emails>**

Hi,

Since the intended disclosure process doesn't look to be working (ticket 2400), here's the disclosure:

Scraped emails from the website, disclosed to them all

Example – The Forge

Tangent – OSINT for contact points

Finding Contact Details ¶

The first step in reporting a vulnerability is finding the appropriate person to report it to. Although some organizations have clearly published disclosure policies, many do not, so it can be difficult to find the correct place to report the issue.

Where there is no clear disclosure policy, the following areas may provide contact details:

- Bug bounty programs such as BugCrowd, HackerOne, huntr.dev, Open Bug Bounty or Standoff.
- A security.txt file on the website at `/.well-known/security.txt` as per RFC 9116. At minimum include the required Contact and Expires fields; optional fields such as Preferred-Languages, Canonical, Policy, and Hiring improve routing for ethical researchers.
- An existing issue tracking system.
- Generic email addresses such as `security@` or `abuse@`.
- The generic "Contact Us" page on the website.
- Social media platforms.
- Phoning the organization.
- Reaching out to the community.

Example – The Forge

10 Jan, Report Received

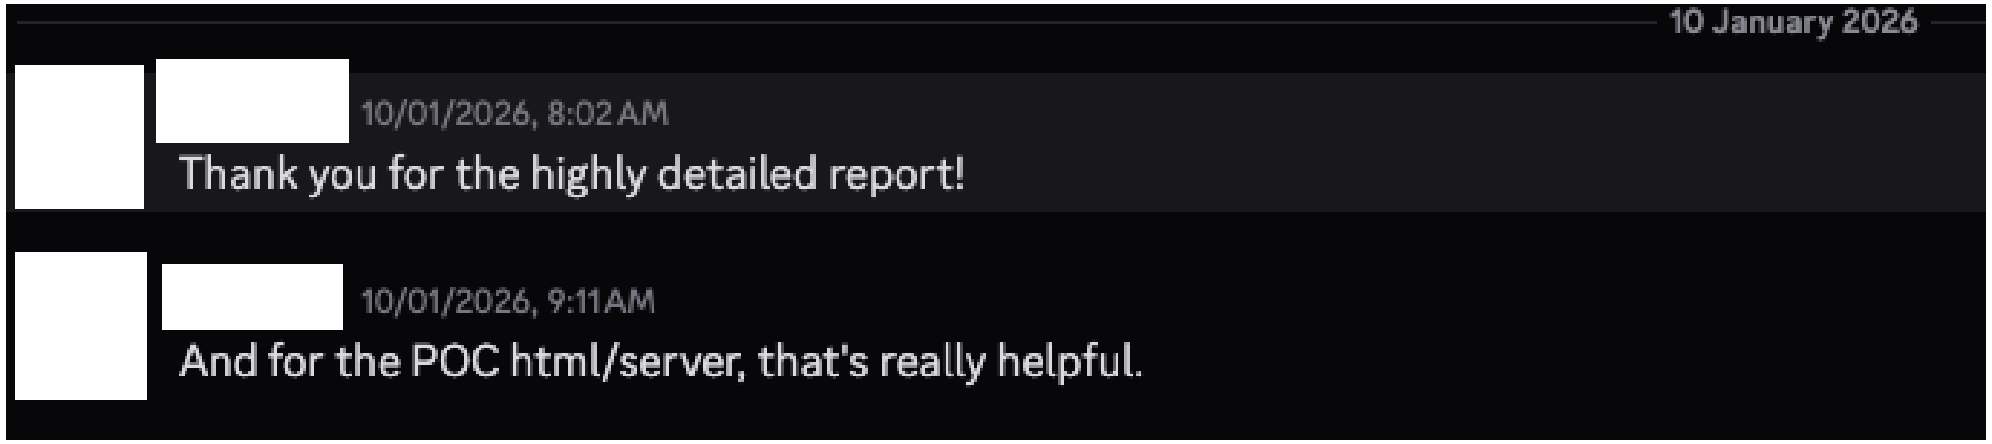

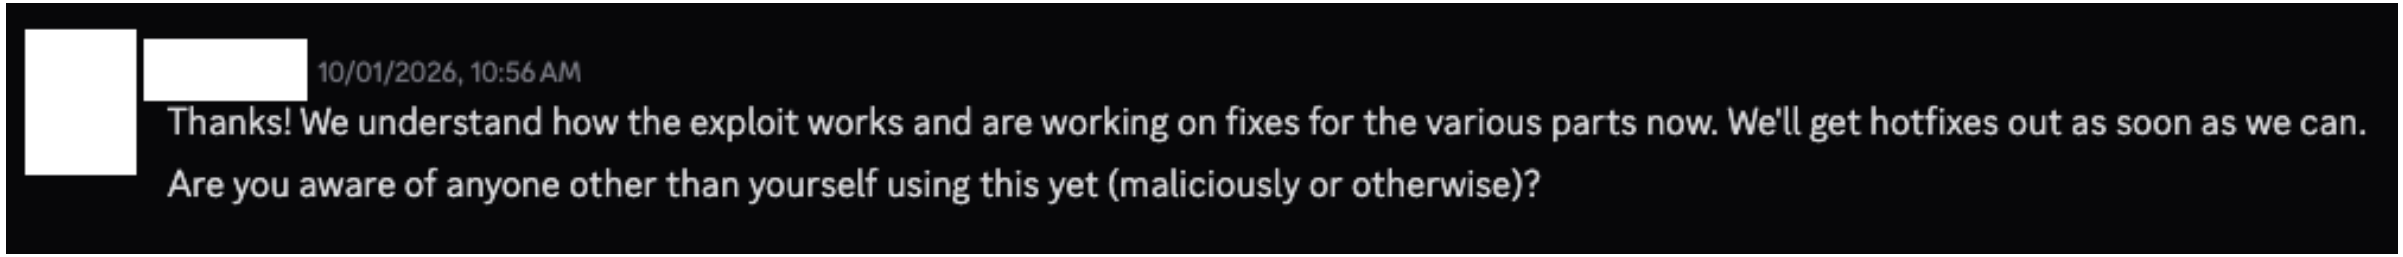

Example – The Forge

14 Jan, Issue Fixed

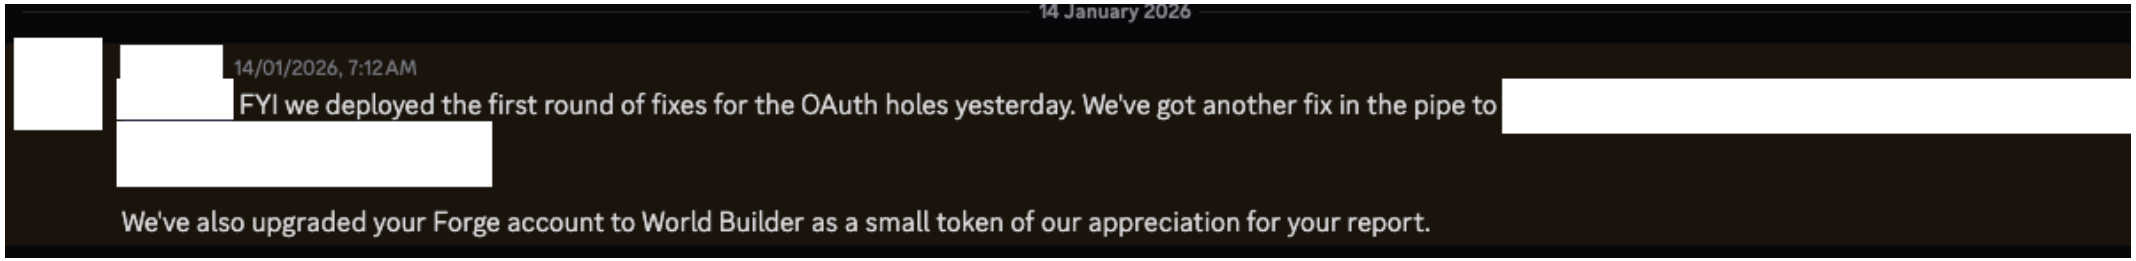

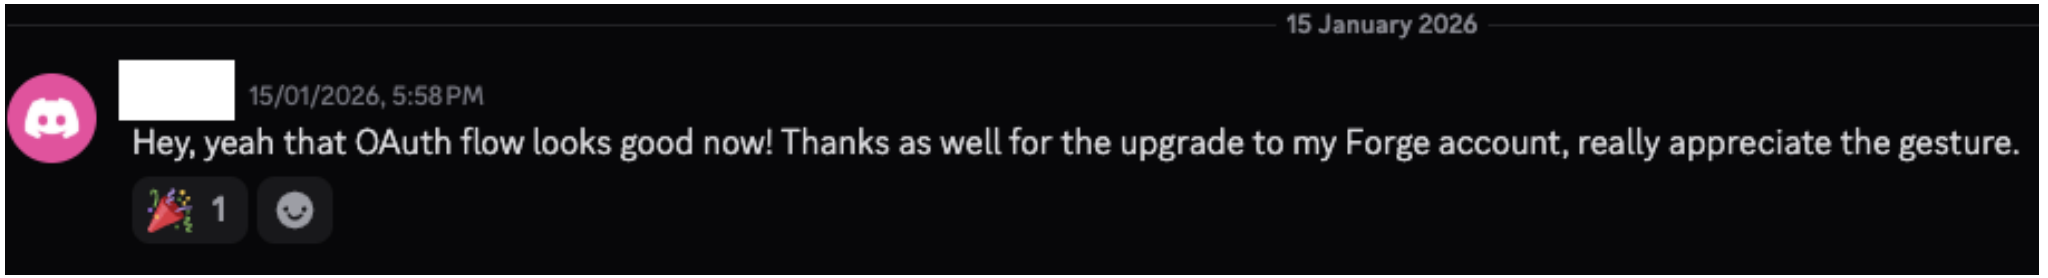

Example – The Forge

Timeline

2 Dec – Asked for security.txt

13 Dec – Permission to test

13 Dec – Asked where to send report

23 Dec – Reporting channel provided

30 Dec – Attempt to disclose

9 Jan – Everyone gets a disclosure!

→ Removed with a security.txt

10 Jan – Triaged

14 Jan - Fixed

→ Quickly fixed once in front of the right person

Example – The Forge

Fallout

- CEO got in contact directly to say thanks personally.
- They gave me free access to their service.
- They're getting their security policy sorted internally.
- Oh, and a job offer!

Is this your email address that you're using to access your The Forge account? If not, provide me which one is in use and I'll get it upgraded to our highest subscription level.

I would also like to know if you are looking for work or a change in job, and if so, would you consider The Forge as your potential employer?

Example – The Forge

What was the bug? No click account takeover via CSRF when visiting a malicious page on a valid subdomain. Specifically:

- No randomised state on their SSO flow when linking an existing account to a SSO provider.
- Arbitrary HTML being able to be hosted on a subdomain of the main site and shared publicly.
- Bunch of smaller issues meaning that it was impossible to invalidate the attacker's session from the victim's account.

Example – The Forge

`state`	**Recommended** Specifies any string value that your application uses to maintain state between your authorization request and the authorization server's response. The server returns the exact value that you send as a `name=value` pair in the URL query component (`?`) of the `redirect_uri` after the user consents to or denies your application's access request. You can use this parameter for several purposes, such as directing the user to the correct resource in your application, sending nonces, and mitigating cross-site request forgery. Since your `redirect_uri` can be guessed, using a **`state`** value can increase your assurance that an incoming connection is the result of an authentication request. If you generate a random string or encode the hash of a cookie or another value that captures the client's state, you can validate the response to additionally ensure that the request and response originated in the same browser, providing protection against attacks such as cross-site request forgery. See the OpenID Connect documentation for an example of how to create and confirm a **`state`** token. ★ **Important:** The OAuth client must prevent CSRF as called out in the OAuth2 Specification . One way to achieve this is by using the **`state`** parameter to maintain state between your authorization request and the authorization server's response.

Example – The Forge

What went wrong?

- No security.txt, had to reach out and ask for permission to test, could have just walked away.
- Long journey to get in front of the right person.
- Unlucky timing clash with an email provider migration.
- Holidays!

Example – The Forge

What went right?

- They offered safe harbour.
- Once it got in front of the right people, it was quickly fixed.
- Recognised that it wasn't a smooth process and the effort that had been put in.
- Let me give this talk to help build up the community!

Example – The Forge

Would I Report Again?

- Despite not having financial reward, yes.
- I've got contacts there now so can avoid the initial difficulties.
- They owned their mistake and made good on it.
- They still need to add a security.txt for others though.

Takeaways

Takeaways

- Consider your process for how vulnerabilities are received.
- Put something in /.well-known/security.txt and define a VDP.
- State your policy around rewards, safe harbour, timelines.
- Get reports in front of the right people but keep an eye on quantities and balancing filtering out garbage to not overwhelm staff.
- Have backup staff to still be able to complete the process when reports come in at inconvenient times.

Links!

OWASP Vuln Disclosure Cheat Sheet

https://disclose.io

Thanks!

Thanks for listening!

My LinkedIn

me@mattdekker.com

Any questions?