Vulnerability Disclosure Policy

1. Introduction

We welcome reports from security researchers who find vulnerabilities in Subcontext. This policy explains what is in scope, how to report a vulnerability, and what you can expect from us. Our contact details are also published at `/.well-known/security.txt`.

Subcontext is used by UK insurers, brokers and other regulated firms to talk to their own customers, so the security of the platform matters to many people beyond us. We would much rather hear about a problem from you than find it exploited. This policy is also published so that our customers can see how external reports are received and handled.

2. Definitions

Term	Meaning
Vulnerability	A weakness in Subcontext's systems that could be used to compromise the confidentiality, integrity or availability of the service or the data in it
Researcher	Anyone who reports a vulnerability to us, whether a professional, a customer or a member of the public
Good faith	Research carried out to find and report vulnerabilities, without intent to cause harm, and in line with this policy

3. Scope

In scope:

- subcontext.com, including the dashboard
- Customer mini-sites hosted at `*.sites.subcontext.com` (the platform, not the content customers publish)
- js.subcontext.com and cdn.subcontext.com
- The Subcontext REST API, MCP server and embeddable widgets
- The Subcontext iOS, Android and Mac apps

Out of scope:

- Customers' own websites and systems, and content customers publish
- Third-party services we use, such as AWS, Stripe and Twilio: please report to them directly
- Denial of service, and high-volume automated scanning
- Social engineering, phishing and physical attacks
- Reports without a demonstrated security impact, such as missing best-practice headers or self-XSS

3.1 Examples of issues we are especially interested in

- Access to another organisation's data, or breaking the separation between customer organisations
- Bypassing authentication, two-factor authentication or role permissions
- Injection, remote code execution or server-side request forgery
- Cross-site scripting that affects other users
- Making an AI agent reveal data it should not, or take an action it should not, in a way that defeats the platform's guardrails or approval steps
- Forging callbacks or webhooks that the platform accepts as genuine
- Weaknesses in the OAuth flows used by the API, MCP server and apps

If you are unsure whether something is in scope, report it anyway and we will tell you.

4. How to Report

Email hello@subcontext.com with the subject "Security vulnerability". Please include:

- The affected URL, app or API endpoint
- A description of the vulnerability and its impact
- Steps to reproduce it
- Your name or handle, if you would like to be credited

Helpful extras include the time of your testing, the account or organisation you tested with, request and response samples, and screenshots or a short video. Please do not include personal data or other customers' data in your report beyond what is needed to show the issue.

5. Rules for Testing

- Only use accounts you own, or have permission to test with
- Access only the minimum data needed to show the vulnerability; if you encounter personal data or another customer's data, stop and report it straight away
- Do not modify or delete data, or degrade the service for anyone else
- Keep the vulnerability confidential until it is fixed, or for 90 days from your report, whichever comes first, unless we agree otherwise

In addition:

- Do not send messages, calls or emails to real people through the platform as part of your testing
- Do not attempt to make payments with cards or bank details that are not yours
- Delete any data you obtained during testing once you have reported it, and do not keep or share it
- Be aware that the platform applies rate limits and a web application firewall, and that sign-in attempts lock an account temporarily. Hitting these is expected, and is not by itself a vulnerability

6. What You Can Expect From Us

- Acknowledgement of your report within 2 business days
- An initial assessment within 5 business days

- Updates as we work on a fix, and confirmation when it is resolved
- Credit for your discovery, if you would like it

7. How We Handle Reports

1. **Receive:** your report is logged and acknowledged.
2. **Triage:** the Security Officer checks the report, reproduces the issue where possible, and assesses its severity and whether any data may have been exposed.
3. **Respond:** if the issue suggests data has been, or may have been, accessed without authorisation, it is handled straight away as a security incident under the Incident Response Plan, including notifying affected customers where required.
4. **Fix:** the fix is developed and released through our normal change process, with priority set by severity in line with the Vulnerability Management Policy.
5. **Confirm:** we tell you when the issue is fixed, and invite you to check the fix if you would like to.
6. **Learn:** we look at why the issue happened and whether similar issues exist elsewhere in the platform.

Worked example

A researcher finds that changing an identifier in an API request returns a lead from another organisation. They stop, and email us the request they used and the one record they saw. We acknowledge within 2 business days, confirm the issue, and treat it as a security incident: we fix the access check, review the logs to establish whether anyone else used the weakness, and notify any affected customer. We then confirm the fix to the researcher and, with their agreement, credit them.

8. Safe Harbour

If you act in good faith and in line with this policy, we will not take legal action against you or report you to law enforcement for your research. If a third party takes action against you over research that followed this policy, we will make it known that you acted with our authorisation.

We consider research carried out in line with this policy to be authorised for the purposes of the Computer Misuse Act 1990. Safe harbour does not apply to activity that is clearly outside this policy, such as extortion, deliberately accessing or keeping customers' data, or disrupting the service.

9. Disclosure

We support coordinated disclosure. Once an issue is fixed, you may publish details of it, and we ask that you share a draft with us first so we can check it does not expose customer data or an unfixed weakness. If a fix will take longer than 90 days, we will explain why and agree a timetable with you.

10. Rewards

We do not currently run a paid bug bounty programme.

11. Related Documents

- Vulnerability Management Policy
- Incident Response Plan

- Secure Development Policy
- Information Security Policy

12. Document Control

Item	Detail
Owner	Security Officer
Approved by	The directors of Subcontext Ltd
Review	At least annually
Effective	October 2026

Subcontext Ltd, registered in England and Wales, Company No. 17019645. Registered office: 124 City Road, London, EC1V 2NX. The current version of this document is published at subcontext.com.