

Vulnerability Disclosure Policy

How to report a security vulnerability in Quorval, and what we commit to in return · Version 1.0 · August 2026

Status: Active policy. Safe-harbour applies to good-faith research conducted within the rules stated below.

Our commitment

If you find a security vulnerability in Quorval, we want to hear about it and we will not pursue legal action against you for reporting it in good faith under this policy. Security research conducted within the rules below is authorised, and we will treat you as a contributor rather than an adversary.

How to report

Email security@quorval.com. Include the affected URL or endpoint, the steps to reproduce, what you were able to access or do, and any proof-of-concept you are willing to share. If the finding is sensitive, say so and we will arrange an encrypted channel. A machine-readable version of these contact details is published at <https://quorval.com/.well-known/security.txt>.

What you can expect from us

Stage	Our commitment
Acknowledgement	Within 2 business days of your report.
Triage and severity assessment	Within 5 business days, with our assessment shared with you.
Remediation target	Critical: 7 days. High: 30 days. Medium: 90 days. Lower severity: next scheduled release.
Status updates	At least every 14 days until the issue is closed.
Credit	Public acknowledgement if you want it, and omission if you do not.

Scope

In scope: the Quorval web application at quorval.com and its subpaths, organization-branded entry paths, our public API endpoints, and our backend functions. Out of scope: infrastructure operated by our providers (report those to the provider), third-party services we integrate with, and any tenant's own content.

Rules of engagement

Do not access, modify, or retain data belonging to any account other than one you control — if you demonstrate a cross-tenant issue, stop at the minimum proof and tell us. Do not run denial-of-service, volumetric, or automated scanning that degrades service for others. Do not use social engineering, phishing, or physical attacks against our staff or providers. Do not publicly disclose before we have remediated or 90 days have elapsed, whichever is sooner, and talk to us if you believe that timeline is wrong for the finding.

Findings we generally consider low value

Missing security headers with no demonstrated impact, reports produced solely by an automated scanner without validation, self-XSS, missing rate limits on non-sensitive endpoints, email spoofing findings on domains with no mail configuration, and version-disclosure banners. We will still read them; they are unlikely to be prioritised.

Rewards

We do not currently operate a paid bug bounty. We are a small company and we would rather tell you that plainly than imply a reward that does not exist. We offer credit, a direct line to the people who fix the issue, and a written confirmation of your finding and its resolution if that is useful to you professionally.

Why this policy exists

A published vulnerability disclosure policy is itself a control that assessors look for, and it is the honest interim answer while an independent penetration test remains outstanding. We would rather have a functioning reporting channel and say the test has not happened yet than have neither and imply otherwise.