Limbus AI Vulnerability Disclosure Policy

Limbus AI is committed to ensuring the security of the system and protecting their information from unwarranted disclosure. Limbus AI has adopted this vulnerability disclosure policy to inform you about our safeguards.

This policy applies to any vulnerabilities you are considering reporting to us. We recommend reading this vulnerability disclosure policy fully before you report a vulnerability and always acting in compliance with it.

Reporting a Vulnerability

If you believe you have found a security vulnerability, please submit your report to us using the following link/email: support@limbus.ai.

Reports may be submitted anonymously.

In your report please include details of:

- Where the vulnerability can be observed.
- A brief description of the type of vulnerability and the title of vulnerability.
- Vulnerability impact.
- Steps to reproduce the observations. These should be a benign, non-destructive, proof of concept. This helps us ensure that the report can be triaged quickly and accurately. It also reduces the likelihood of duplicate reports.

What to Expect

After you have submitted your report, we will respond to your report within 2 working days and aim to triage your report within 10 working days. We'll also aim to keep you informed of our progress.

Priority for remediation is assessed by looking at the impact, severity, and exploit complexity. Vulnerability reports might take some time to triage or address.

We will notify you when the reported vulnerability is remediated, and you may be invited to confirm that the solution covers the vulnerability adequately.

Rules of Engagement

You must **not**:

- Break any applicable law or regulations.
- Access unnecessary, excessive or significant amounts of data.
- Modify data in Limbus AI systems.
- Use high-intensity invasive or destructive scanning tools to find vulnerabilities.
- Attempt or report any form of denial of service, e.g. overwhelming a service with a high volume of requests.

- Disrupt Limbus AI services or systems.
- Submit reports detailing non-exploitable vulnerabilities, or reports indicating that the services do not fully align with “best practice”.
- Communicate any vulnerabilities or associated details other than by means described in this policy.
- Social engineer, ‘phish’ or physically attack Limbus AI staff or infrastructure.
- Demand financial compensation in order to disclose any vulnerabilities.

You must:

- Always comply with data protection rules and must not violate the privacy of Limbus AI users, staff, contractors, services or systems.
- Securely delete all data retrieved during your research as soon as it is no longer required or within 1 month of the vulnerability being resolved, whichever occurs first, or as otherwise required by Limbus AI Privacy Policy.

Disclosure

Limbus AI is committed to timely correction of vulnerabilities. However, we recognize that public disclosure of a vulnerability in absence of a readily available corrective action likely increases risk of exploitability. Limbus AI will not publicly disclose or discuss any potential vulnerability until it has been confirmed and, if applicable, a software patch is available. Accordingly, we require that you refrain from sharing information about discovered vulnerabilities. If you believe others should be informed of the vulnerability prior to our implementation of corrective actions, we require that you coordinate in advance with us.

Limbus AI may share vulnerability reports with the appropriate Security Agencies, as well as any affected vendors. We will not share names or contact data unless given explicit permission.

Legalities

This policy is designed to be compliant with common vulnerability disclosure good practice. It does not give you permission to act in any manner that is inconsistent with the law, or which might cause Limbus AI or any partner organization to be in breach of any legal obligations.