LatentWorlds AI

THE DEPLOYMENT BACKEND FOR PHYSICAL AI

Public | Version 1.0 | 2026

Security and Responsible Disclosure Policy

Published 2026-02-24

> **Purpose statement:** This policy describes our security principles and how to report potential vulnerabilities in LatentWorlds AI systems.

Document status: Public
Version: 1.0
Effective date: 2026-02-24
Owner (top management): Alejandro Daniel Noel, Cofounder & CTO
Website publication location: https://www.latentworlds.ai/policies/security
Contact: policy@latentworlds.ai

1. Security principles

LatentWorlds builds infrastructure for robotics and physical AI. Reliability and security are core product requirements. Our security program is sized to our stage, and we continuously improve it as we scale.

Our security principles include:

- **Least privilege by default.** Access to production systems is limited to the minimum necessary and reviewed as roles change.
- **Defense in depth.** We use multiple layers of controls (identity, network controls, encryption, logging) to reduce the impact of any single failure.
- **Auditability.** We design for traceability of sensitive actions and changes.
- **Secure development.** We use code review, dependency hygiene, and automated checks where feasible.
- **Prompt remediation.** We prioritize fixes based on risk and potential impact.

2. Reporting security issues

If you believe you have found a security vulnerability in our Website or Services, please report it as soon as possible by emailing:

`policy@latentworlds.ai`

Include “Security” in the subject line and provide:

- a clear description of the issue and affected component(s);
- steps to reproduce (proof-of-concept where helpful);
- potential impact and any assumptions; and
- your preferred contact details for follow-up.

If you need to share sensitive details, you can request an encrypted channel in your initial email.

3. Guidelines for responsible research

We welcome good-faith security research. When conducting research, you agree to:

- avoid privacy violations, data destruction, service disruption, or social engineering;
- use only the minimum amount of testing needed to confirm the vulnerability;
- stop testing and report promptly if you encounter personal data or confidential customer data; and
- not publicly disclose details until we have had a reasonable opportunity to investigate and remediate.

This policy does not authorize access to data or systems beyond what is necessary to identify and report vulnerabilities.

4. Safe harbor

We will not pursue legal action against you for good-faith security research that:

- is consistent with this policy;
- is aimed at improving security and protecting users; and
- does not involve extortion, theft, or attempts to access or modify data beyond what is necessary to demonstrate the issue.

If you are uncertain whether your planned testing is within scope, contact us first.

5. Our disclosure process

When we receive a report, we aim to:

- acknowledge receipt within **5 business days**;
- provide a status update within **15 business days**; and
- remediate as quickly as reasonably possible, based on severity and operational constraints.

We may request additional information, coordinate on timelines, and—where appropriate—credit reporters in release notes or acknowledgements (with your consent).

We do not currently operate a bug bounty program.

6. Scope

This policy covers:

- the Website at https://www.latentworlds.ai; and
- LatentWorlds-hosted Services and related APIs, as applicable.

Customer-managed deployments may have additional constraints and reporting requirements. If you are testing in a customer environment, obtain the customer's explicit permission and follow the customer's policies.

7. Changes to this policy

We may update this policy from time to time. We will publish the updated version on our website and update the effective date.

8. Adoption and signature

By signing this document, top management adopts this Security and Responsible Disclosure Policy for publication and use.

Adopted on: 2026-02-24

Alejandro Daniel Noel
Cofounder & CTO, LatentWorlds AI

Signature: