

Acceptable Use Policy

Anweshna AI Due Diligence Platform

Version 1.1 — Effective 28 September 2026

Downloaded from <https://anweshna.com/aup.html> · the web page is the authoritative version.

This AUP forms part of the [Terms of Service](#); capitalised terms have the meanings given there. Violations may lead to suspension or termination under Terms §13.

1. Prohibited Content

You must not upload, store, or process through the Platform any content that:

- you lack the legal right to process or share with us and our sub-processors (including material subject to legal professional privilege, protective orders, or confidentiality undertakings that do not permit third-party processing);
- is unlawful in your jurisdiction or ours, including stolen or hacked data and unlawfully intercepted communications;
- contains child sexual abuse material, terrorism content, or content facilitating serious crime — **zero tolerance; reported to authorities where required**;
- contains malware, exploits, or executable payloads disguised as documents;
- comprises special-category or highly regulated data (health records, biometrics, payment-card databases, government ID databases, classified material) without our prior written agreement;
- is uploaded to build a profile of a private individual unconnected to a legitimate business transaction.

2. Prohibited Conduct

- Accessing or attempting to access another client's data, deal rooms, or reports;
- Probing, scanning, penetration-testing, load-testing, or otherwise testing the security or capacity of the Platform without our prior written authorisation;
- Circumventing authentication, rate limits, seat limits, plan gating, deal-room expiry, clean-room or clean-team restrictions, Release Airlock review, or encryption;
- Sharing API keys or sub-keys outside your organisation, or pooling a single account across unrelated organisations;
- Scraping the Platform, harvesting data about other users, or submitting forged or spoofed requests (including forged telemetry, lead, or consent submissions);
- Reverse-engineering, decompiling, or extracting the Platform's prompts, rulesets, or scoring logic, including via adversarial prompting of the analysis engines ("prompt injection" probing);
- Uploading documents crafted to manipulate the AI's output — e.g. embedded instructions intended to suppress risk findings or alter scores — where the resulting Output will be shown to any other party;
- Reselling, sublicensing, or operating the Service for third parties without a written agreement with us;
- Using the Service or Output to build, train, or benchmark a competing product.

3. Legal & Regulatory Misuse

- **Securities laws:** no use in furtherance of insider trading, tipping, or market manipulation. Documents and Output relating to non-public transactions are MNPI — handle them under your information-barrier policies.
- **Sanctions & export controls:** no use from, for, or on behalf of comprehensively sanctioned jurisdictions or restricted parties (Terms §17).
- **Sole-basis decisions:** Output must not be the sole basis for investment decisions, regulatory filings, court submissions, or any assessment of a natural person's eligibility for employment, credit, insurance, or housing.
- **Impersonation:** no signing up or accepting terms on behalf of an entity you are not authorised to bind.

4. API & Resource Usage

- Respect published rate limits and plan allowances; no artificial request amplification or parallel-key evasion of limits;
- No automated re-analysis loops of unchanged documents to farm output variations;

- Public endpoints (signup, contact, telemetry, feedback) may only be called for their intended purpose from our own pages or documented flows.

4.1 Autonomous agents (MCP)

An AI agent connected through our MCP endpoint acts under your API key and is subject to everything in this policy. In addition, you must not:

- Run an agent in an unattended loop that re-screens documents, polls job status, or retries failed calls without a bounded stop condition;
- Point an agent at the Service to enumerate, scrape, or mirror content, or to reconstruct our rulesets, prompts, or scoring behaviour;
- Grant an agent broader access to your account than the task requires — use a role-scoped sub-key where the work does not need owner rights;
- Let an agent act on Output without the human review required by Section 11.2 of the [Terms](#), or use it to take an action this policy prohibits a person from taking.

Documents may contain text intended to manipulate an AI system that reads them, and findings are returned into your agent's context. Deliberately uploading a Document crafted to make an agent — yours or another client's — behave in a way this policy prohibits is itself a violation, and is covered by the prohibition on score manipulation and prompt-injection probing in Section 2.

5. Security Research

We welcome responsible disclosure. If you believe you have found a vulnerability, report it to security@anweshna.com and do not access other clients' data, exfiltrate more than needed for a proof of concept, or disclose publicly before we have had a reasonable opportunity to remediate. Good-faith research within these bounds will not be treated as an AUP violation; we will not pursue legal action against good-faith researchers who comply with this section.

6. Enforcement

We may investigate suspected violations, preserve and review relevant logs, throttle or suspend access, remove offending content, terminate for material breach, and notify authorities where legally required. Where practical and lawful, we will notify you and give you a chance to cure before suspension. We are not obligated to monitor content, and no failure to enforce is a waiver.

7. Reporting

Report abuse, suspected account compromise, or AUP violations to abuse@anweshna.com.

Anweshna AI Due Diligence · E46 Mahadevpur 4, Namsai-792105, Arunachal Pradesh, India

Legal: legal@anweshna.com · Privacy: privacy@anweshna.com · General: info@anweshna.com

Anweshna scores documents for risk signals. It does not replace legal review, independent verification, or expert due diligence.