

Terms of Service & Acceptable Use

Version 16 September 2026. Applies upon affirmative acceptance.

The agreement

These Terms of Service & Acceptable Use (the “Terms”) are an agreement between Adverserial AI LLC, a New Jersey limited liability company, operating under the brand Adverserial AI (“Company,” “we,” “us”), and the person or organization accepting them (“Customer,” “you”). They govern our hosted cybersecurity AI services, including CyberKimi, CyberGLM, chat interfaces, APIs and associated features (the “Services”). Separately licensed software and model weights remain subject to their applicable licenses.

“Input” means content you submit or make available to the Services. “Output” means content generated in response. Together they are “Customer Content.” “Users” means people you authorize to use your account.

These Terms become binding when you affirmatively accept them through an account, purchase or other acceptance process that clearly presents them, or execute an agreement incorporating them. Merely viewing this page is not acceptance. If you accept for an organization, you represent that you have authority to bind it. Do not activate or use the Services under these Terms if you do not agree.

A signed enterprise agreement or order expressly addressing a conflicting provision controls that conflict. A data processing agreement (“DPA”) controls conflicting provisions about personal-data processing. An order’s agreed commercial terms control conflicting general billing terms here. Mandatory law always applies. Descriptions of possible enterprise configurations do not, by themselves, create a deployment commitment, security certification or service-level agreement (“SLA”).

Eligibility & account control

You must be at least 18 years old and legally capable of entering this agreement. The Services are designed for professional and other lawful cybersecurity work, not for children. Provide accurate account and billing information and keep it current.

You control your Users, credentials, API keys, permissions and integrations. Apply least-privilege access, protect secrets, remove access when roles change, and promptly report suspected compromise. Do not sell, transfer or share access contrary to your order or evade seat, usage or access restrictions. Organization administrators may manage Users and view organizational usage as described for the service configuration.

You are responsible for activity you authorize and failures to exercise reasonable care over your account. Responsibility for unauthorized activity will be determined under applicable law and the parties’ respective conduct; this provision does not transfer responsibility for a security failure caused by us to you.

Cybersecurity AI & its limitations

The Services can produce sensitive, dual-use cybersecurity material. This may include exploit analysis, vulnerability research, offensive techniques, malware analysis and code that could cause harm if misused. Models intended to support advanced cybersecurity research may respond to technical requests that a general-purpose assistant would restrict.

Do not assume that the Services have guardrails capable of preventing unlawful, unsafe or damaging activity. A response is not a finding that your request is lawful, accurate, authorized or safe. A refusal is not a complete assessment of risk. Technical capability never supplies permission to act.

AI can invent facts, citations, vulnerabilities and dependencies; miss real threats; misunderstand scope; generate insecure or destructive code; and behave differently across repeated requests. Confidence, detail or a polished explanation is not evidence of correctness. Model behavior, performance, context capacity and availability may change, subject to express contractual commitments.

We do not promise that every request will be answered or that safeguards, usage restrictions or legal controls will never apply. We may use restrictions necessary for legal compliance, service integrity or enforcement of these Terms.

Authorization comes first

You must have a lawful basis and all necessary permissions before submitting data, connecting a tool or applying Output. Before testing a third-party system, obtain and maintain legally sufficient authorization, ordinarily documented in writing, identifying the relevant assets, techniques, time windows, data handling and limitations. Confirm that the person granting permission has authority to do so.

Stay within the approved scope, including limits on exploitation, persistence, social engineering, data access and disclosure.

Reconfirm authorization when the target, method, ownership or engagement changes. Stop when permission expires or is withdrawn.

For bug bounties, follow the actual program terms. Public accessibility, a discovered weakness or possession of credentials does not itself authorize access.

For client work, maintain appropriate client agreements and obtain permission for third-party AI processing where required.

Possess and analyze samples, code, credentials and datasets only where you have the necessary rights and lawful purpose.

You remain responsible for determining which laws, professional duties and contractual restrictions apply. We do not grant rights in any target system or certify that your activities qualify as lawful security research.

Permitted research. Prohibited harm.

Lawful security research, authorized red-team and blue-team operations, defensive engineering, reverse engineering, education and controlled testing are permitted, subject to your agreement and applicable law. Offensive terminology or techniques alone do not determine whether work is authorized.

You must not use the Services to conduct, facilitate or knowingly assist:

Unauthorized access, access beyond permission, theft of credentials or data, unlawful exfiltration, or continued access after permission is withdrawn.

Malicious deployment of malware or ransomware, extortion, fraud, credential harvesting, or deceptive phishing outside a lawful, authorized engagement.

Unlawful disruption, destructive attacks, sabotage, or concealment of unauthorized compromise.

Unlawful surveillance, privacy violations, harassment, exploitation of children, trafficking, terrorism or other violent criminal conduct.

Infringement or misappropriation of intellectual property, trade secrets or confidential information.

Evasion of applicable sanctions, export restrictions, legal access restrictions, service security controls, metering or billing.

Abusive automated traffic, extraction of other customers' data, or attacks on the Services without our separate written authorization.

Controlled simulations of attack techniques must be lawful, appropriately contained and specifically authorized. Labeling conduct "research," "education" or "testing" does not make unlawful conduct permissible. You may not direct another person, agent or downstream application to do what you are prohibited from doing yourself.

Human review & execution safeguards

You decide whether and how to use, execute, publish or distribute Output. Before taking action, a competent person must independently review the relevant code, assumptions, targets, evidence and likely consequences. Do not treat generated text as an executable instruction merely because the model suggests it.

Test in an isolated environment before production use; verify dependencies, privileges, network destinations and destructive operations.

Use target allowlists, restricted credentials, approval gates, execution timeouts, rate limits and spending limits appropriate to the risk.

Maintain suitable backups, rollback procedures and incident-response plans.

Treat documents, repositories, web pages and other retrieved content as untrusted. Defend against prompt injection and unauthorized tool use.

Check findings before escalation or publication, and follow applicable disclosure obligations.

Do not rely on AI as the sole basis for decisions where an error could foreseeably cause death, serious injury or major physical harm. Work involving critical infrastructure requires qualified oversight, valid authorization and appropriate controls. The Services are not an emergency-response system or a substitute for competent professional judgment.

Agents, tools & external services

If you enable agents, tool calls, code execution or integrations, actions may modify or delete files, change systems, expose information, send communications or incur charges. Understand the permissions you grant and maintain approval and containment controls. Authorization to use an integration is not authorization to act on every resource it can technically reach.

Third-party services you independently connect are subject to their own terms and data practices. Review them before transferring Customer Content. We are not responsible for independently selected third-party services except to the extent required by law or expressly agreed. This does not excuse our obligations for our own selected subprocessors or subcontractors.

Unless expressly contracted, we provide AI capabilities rather than an engagement to operate your security program, validate your authorization, or supervise each action. Any separately managed service is governed by its agreed scope.

Customer Content & intellectual property

You retain your rights in Input. You represent that you have the rights, permissions and lawful bases necessary to submit it and permit its processing. You grant us a limited license to process Customer Content as necessary to deliver and support the Services, comply with applicable law, and carry out your instructions, subject to applicable privacy commitments and any DPA. This is not a license to sell your content or train models on it.

As between you and Company, to the extent Company has transferable rights in Output, Company assigns those rights to you. This does not transfer underlying models, service software or pre-existing Company materials, and remains subject to third-party rights and applicable law. AI-generated material may not qualify for copyright protection, may resemble other outputs and is not guaranteed to be unique or non-infringing. Assess licenses and rights before use or distribution.

Company and its licensors retain rights in the Services, models, branding and underlying technology. Access grants a limited right to use the Services under your agreement, not ownership of model weights or infrastructure. Do not bypass technical controls or misuse our intellectual property. Nothing here restricts rights that applicable law makes non-waivable.

Privacy, confidentiality & hosting

[Read our Privacy Policy](#)

We do not use Customer Content to train or fine-tune models unless you separately and expressly opt in. Any such permission must identify the agreed purpose and scope. Silence or acceptance of a later general terms update is not permission for a new, incompatible training use.

Hosted inference necessarily processes Input and Output. These Terms do not represent that we receive no data, that no operational metadata exists, or that every deployment has zero retention. The applicable privacy notice, deployment documentation and DPA must specify actual processing, retention, deletion, hosting regions and subprocessors. Obtain those details before submitting information that depends on a particular configuration; contact contact@adversarial.ai if they have not been provided.

Chat history, caching, backups, billing records, security metadata and support attachments can have different lifecycles. A zero-content-retention commitment applies only where expressly agreed and technically configured; it does not automatically erase records we must lawfully retain. Customer-controlled integrations may retain their own copies. Hosting location, private tenancy, customer-cloud deployment and air-gapping are distinct features and require explicit agreement.

We will protect non-public Customer Content using reasonable safeguards appropriate to the Services and restrict its use and disclosure to authorized service purposes and legal obligations. Personnel and provider access must be limited to what is necessary and permitted under the applicable agreement. No system is perfectly secure; this disclosure does not waive our security, confidentiality or data-protection duties.

Do not submit secrets, personal data, regulated information or third-party confidential material unless you have a lawful basis and the selected configuration and agreements support that processing. Where legally required, a DPA must be in place first. Do not assume legal privilege, a particular regulatory certification, or an

industry-specific compliance agreement merely from use. Legally required disclosures remain subject to applicable safeguards and notice restrictions.

Fees, commitments & cancellation

Prices, included usage, any overage rates, renewal periods and deployment costs are those clearly disclosed at checkout or in your order. These Terms do not create an unlimited-use entitlement. Taxes and third-party charges apply only as disclosed or required by law. You authorize charges you affirmatively agree to and must provide valid payment information.

Recurring subscriptions require clear disclosure and your consent before purchase. Cancel future renewal through available account controls or by contacting us; mandatory cancellation rights and procedures remain available. Cancellation ordinarily takes effect at the end of the paid term unless the order or law provides otherwise. We will not apply a price increase retroactively to usage already incurred.

Credit expiry, rollover, minimum commitments, included models, capacity and refund conditions must be disclosed in the purchase terms. We may not introduce undisclosed limits after purchase to defeat agreed usage. Statutory refund rights are preserved. If we permanently discontinue a prepaid service without supplying the contracted alternative, we will refund the unused, undelivered portion unless a lawful signed agreement provides a different remedy.

Promptly raise suspected billing errors with enough information to investigate. No administrative reporting period here extinguishes legal rights. Service credits and prepaid usage balances are contractual entitlements, not bank deposits or investment products.

Availability & service changes

Unless a signed SLA applies, no specific uptime, response time, throughput, context size under every load, or recovery time is guaranteed. Published capabilities may depend on configuration, workload and capacity. Beta, preview and experimental features may be incomplete and must not be treated as production commitments.

We may maintain and improve the Services, address vulnerabilities, and replace or retire features. For a material reduction to a paid commitment, we will provide reasonable notice and an appropriate contractual remedy, except where urgent legal or security needs require immediate action. Express enterprise commitments and mandatory consumer rights remain controlling.

A private deployment does not imply unlimited capacity, geographic failover or an SLA without an agreed design. Performance and isolation commitments must be specified in the applicable order.

Warranty disclaimer

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, AND EXCEPT FOR EXPRESS WARRANTIES IN A SIGNED AGREEMENT, THE SERVICES AND OUTPUT ARE PROVIDED “AS IS” AND “AS AVAILABLE.” COMPANY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT, AND WARRANTIES ARISING FROM COURSE OF DEALING OR USAGE OF TRADE.

We do not warrant that Output will be correct, complete, secure, lawful for your intended use, free of harmful code, or capable of detecting every vulnerability or threat. We do not warrant uninterrupted or error-free operation or a particular commercial, research or security outcome. Independently verify material conclusions.

These disclaimers do not negate an express contractual promise, authorize deceptive claims, or exclude warranties or remedies that cannot lawfully be excluded. If a disclaimer is unenforceable in your jurisdiction, it applies only to the lawful extent.

Allocation of responsibility

You are responsible for your authorization, instructions, selected targets, submitted content, independent review, execution decisions and downstream use. As between the parties, you bear the consequences of your unlawful use, misuse, or failure to implement appropriate safeguards, to the extent caused by that conduct. For organization accounts, this includes conduct of Users and contractors acting under your authority.

You acknowledge the disclosed operational risks of AI-generated cybersecurity material. Deploying Output without appropriate review or safeguards can create serious harm. Neither our generation of Output nor our failure to reject a request is approval, a representation of legality, or an agreement to assume your operational responsibilities.

This allocation does not excuse Company's own unlawful conduct or obligations imposed by law. It does not bind third parties, prevent a regulator from acting, or determine criminal responsibility. Responsibility depends on the facts, applicable law and the enforceable agreements between the parties.

Limitations of liability

TO THE MAXIMUM EXTENT PERMITTED BY LAW, COMPANY AND ITS AFFILIATES, LICENSORS AND SERVICE PROVIDERS WILL NOT BE LIABLE UNDER THIS AGREEMENT FOR INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, EXEMPLARY OR PUNITIVE DAMAGES, INCLUDING CONSEQUENTIAL LOSS OF PROFITS, REVENUE, BUSINESS, GOODWILL, OPPORTUNITY OR DATA, EVEN IF ADVISED OF THEIR POSSIBILITY.

SUBJECT TO THE EXCEPTIONS BELOW, THEIR COMBINED AGGREGATE LIABILITY ARISING OUT OF OR RELATING TO THE SERVICES OR THESE TERMS WILL NOT EXCEED THE GREATER OF US\$100 OR THE FEES YOU PAID COMPANY FOR THE AFFECTED SERVICES DURING THE 12 MONTHS BEFORE THE FIRST EVENT GIVING RISE TO THE CLAIM.

The limits apply across contract, tort (including ordinary negligence), strict liability and other legal theories to the extent lawful. Related events are treated as one series for the cap, and multiple claims do not multiply it. A signed agreement may establish different limits.

Exceptions: nothing in these Terms excludes or limits liability for fraud, fraudulent misrepresentation, willful misconduct, gross negligence, or any liability that applicable law prohibits us from excluding or limiting. This includes non-waivable statutory, consumer and data-protection rights, and death or personal injury where limitation is prohibited. No provision limits enforcement powers, criminal liability, or non-waivable rights of a person who is not bound by these Terms.

These provisions allocate commercial risk and apply only to the extent enforceable. They do not promise that all damage connected with the Services is legally the Customer's responsibility. If a limitation is unavailable, the remaining lawful provisions continue to apply.

Business-customer indemnification

If you use the Services for a business or organization, you will defend Company and its personnel against third-party claims, and indemnify them for damages finally awarded, reasonable legal costs, and settlements you approve, to the extent caused by your or your authorized Users' unlawful use of the Services, material breach of the authorization or acceptable-use obligations, or submission of Input that infringes a third party's rights.

This obligation does not cover the portion caused by Company's breach, negligence or other wrongful conduct, or liability that cannot lawfully be indemnified. It does not transfer criminal responsibility or require payment of a governmental penalty where reimbursement is prohibited. It is not a consumer indemnity for ordinary personal use.

We must provide reasonably prompt notice (delay reduces the obligation only to the extent it materially prejudices the defense), reasonable cooperation at your expense, and reasonable control of the defense using qualified counsel. We may participate at our own expense. You may not settle a claim in a way that admits our fault, imposes non-monetary obligations on us, or fails to release us without our prior written consent, not to be unreasonably withheld.

Suspension, termination & enforcement

We may restrict or suspend access where reasonably necessary to address suspected unlawful activity, a material breach, non-payment, a security threat or a binding legal requirement. We will use proportionate measures and provide notice and an opportunity to resolve the issue where practical, unless doing so would create risk or violate law. A serious or repeated breach may lead to termination.

Enforcement may use reports and information lawfully available to us under applicable privacy terms and agreements. This section does not permit undisclosed content retention or monitoring. Nor does it promise that we review every request or detect every misuse. We may preserve or disclose information where legally required, subject to applicable obligations.

You may stop using the Services and close your account, subject to accrued fees and lawful commitments. Upon termination, access rights end; data handling follows the applicable retention schedule and law. We will provide any export opportunity or refunds required by the agreement or law. Contact us to challenge an access restriction or report an error. Provisions intended to survive, including accrued payment obligations, content rights, confidentiality, lawful liability limits and dispute provisions, survive termination.

Abuse, security & copyright reports

Send concerns to contact@adverserial.ai. Include a reply address, relevant account or request identifiers, a description and supporting information you can lawfully share. Do not email passwords, live credentials or unnecessary personal data. Ask for an appropriate secure channel when sensitive evidence is involved.

For an intellectual-property report, identify the protected work or right, the allegedly infringing material and its location, your authority to report, and why you believe the use is unauthorized. We may request additional information and take appropriate measures under applicable law, including restrictions for repeat infringement where required.

This general contact address is not a representation that a statutory DMCA agent registration has been completed. Any formal statutory notice procedure will use the legally required contact and process applicable to the service concerned. We do not claim a copyright safe harbor merely by publishing these Terms.

Legal compliance & U.S. framework

You must comply with applicable computer-misuse, privacy, intellectual-property, anti-fraud, sanctions and export-control laws. Do not use or supply the Services to prohibited persons, locations or end uses, or misrepresent identity or location to evade restrictions. Restrictions depend on the transaction and law; these Terms do not classify any model or authorize an export.

Company reserves defenses and statutory protections only where they actually apply. The following references explain important boundaries; none is a blanket immunity for cybersecurity AI :

47 U.S.C. § 230 □ Provides certain protections concerning information supplied by another content provider, with statutory exceptions. Its application to AI-generated material is fact-dependent; it must not be assumed to protect content the Company creates or develops.

17 U.S.C. § 512 □ Offers conditional copyright safe harbors for qualifying activities. Eligibility and required procedures matter; these Terms do not establish that the Company meets them.

15 U.S.C. § 7001 — electronic contracting □ Electronic form alone does not invalidate a contract. This does not eliminate valid-assent requirements, applicable consumer disclosures or substantive legal limits.

18 U.S.C. § 1030 — computer misuse □ Governs specified computer-related offenses and civil claims. A service subscription is not authorization to access another person's computer. Government charging policies are not permission or a general defense to civil claims.

18 U.S.C. § 2 — aiding and abetting □ A private contract does not excuse a party's own criminal participation. Allocating customer responsibilities cannot transfer or eliminate criminal liability.

15 U.S.C. § 45 — unfair or deceptive practices □ Disclaimers do not permit deceptive statements about capabilities, privacy or security. Applicable statutory duties continue to apply.

California Civil Code § 1668 □ An example of state-law restrictions on contractual exemptions for fraud, willful injury or violations of law. Other jurisdictions impose their own limits. This reference is not a selection of California law.

Disputes & applicable law

If a dispute arises, please contact us with a description and the remedy sought. The parties will try in good faith to resolve it informally for 30 days where practical. This process does not prevent urgent relief, a regulatory complaint, or a filing needed to preserve a limitation period or mandatory right.

If a signed agreement specifies governing law and venue, that provision applies to the extent enforceable. Otherwise, applicable law and jurisdiction are determined by a court of competent jurisdiction under its applicable rules, including mandatory protections available to you.

These Terms do not impose mandatory arbitration, a class-action waiver or a jury-trial waiver. Nothing restricts lawful reporting to authorities or cooperation with an investigation.

Changes to these Terms

We may update these Terms prospectively. We will identify the updated version and provide at least 30 days' advance notice of material changes through the account or registered contact, unless a shorter period is necessary for legal compliance or an urgent security issue. Such exceptions will be limited to what is necessary.

We will obtain new affirmative agreement where required by law or where the change requires fresh consent. We will not use a general update to authorize an incompatible new use of previously submitted Customer Content. Changes do not retroactively rewrite accrued claims, incurred charges or a signed agreement. If you do not accept an applicable change, you may stop using the affected Services and exercise any contractual or statutory cancellation rights.

General terms & notices

These Terms, incorporated purchase terms and applicable signed agreements form the agreement about the Services. Neither party is the other's agent, partner or employer. A failure to enforce a provision is not a waiver of later enforcement. If a provision is invalid, it is limited or severed only as law permits; the remaining enforceable agreement continues.

You may not assign the agreement without our consent, not to be unreasonably withheld. Company may assign it as part of a genuine merger, reorganization or transfer of the relevant business, provided the successor assumes the obligations and the transfer respects privacy commitments and applicable law. Mandatory rights to object or terminate remain available.

Neither party is liable for delay caused by an event genuinely beyond its reasonable control to the extent the law permits, provided it takes reasonable mitigation steps. This does not excuse reasonable security measures, confidentiality duties, payment for services already delivered, or mandatory remedies for undelivered services.

We may send operational and contractual notices to your registered contact or account. Maintain a valid contact address. Where electronic delivery requires specific consent or another form of notice is legally required, that requirement controls. Our licensors and service providers may rely on the lawful protections expressly extended to them, but receive no greater protection than applicable law permits.

Contact & agreement record

Adverserial AI LLC Operating as Adverserial AI 1078 Summit Ave #605 Jersey City, NJ 07307, USA
contact@adverserial.ai

For terms, abuse reports, privacy documentation, billing concerns or enterprise contract requests, use the contact details above. Email does not replace a method of legal service required by law.

Keep a copy of the version you accept. The account or order acceptance process, not a visit to this page, records your agreement. A signed enterprise agreement may supplement these Terms with the parties' legal addresses, hosting region, data-retention schedule, security commitments, SLA and negotiated liability allocation.