BANKR SECURITY

Bug Bounty & Vulnerability Policy

Scope, severity classification, rewards, and disclosure terms

Platform: bankr.bot
Version: 1.0
Effective: July 2026
Contact: security@bankr.bot

1 Purpose & Principles

This Policy describes how Bankr invites and handles good-faith security research on Bankr-controlled systems, including what is in scope, how to report vulnerabilities, what we offer by way of rewards, and how we protect researchers who follow the rules. It supplements Bankr's internal vulnerability handling and incident-response processes and is not a substitute for those internal procedures.

We welcome good-faith security research across the Bankr platform; §2 defines the specific assets in and out of scope. Three principles govern how we read every report:

- **Primacy of impact.** Severity is driven by demonstrated impact — how many users are affected and whether funds are at risk — not by the theoretical class of the bug.
- **Primacy of rules.** Safe harbor and rewards apply only to research conducted within the Rules of Engagement (§8). Impact obtained by breaking those rules does not qualify.
- **Reproducibility.** A report must show a concrete, working exploitation path, not a scanner finding or a hypothetical.

Two dimensions dominate our valuation of any report: **blast radius** (one user vs. all users at once) and **financial impact** (loss, theft, or freezing of funds). Both drive the categories and severity model in §4.

This Policy is intended as a public statement of Bankr's approach to vulnerability disclosure and bug bounties. It does not by itself create a contract, offer, or promise of payment, and Bankr may choose, in its discretion, whether and how to apply it in particular circumstances.

2 Scope

Test only assets Bankr owns and operates. Scope is intentionally high-level and centers on the Bankr terminal at **bankr.bot** and the surfaces it exposes; if unsure whether something is in scope, ask at security@bankr.bot or Bankr Help Center before testing.

Some in-scope systems may rely on third-party infrastructure or components that are subject to separate contractual, licensing, or legal restrictions. Where a reported issue primarily affects an upstream vendor or service provider, Bankr may need to coordinate with that provider and may be limited in the information it can share about their systems or remediation steps. This Policy does not endorse, authorize, solicit, or reward any unauthorized access to third-party infrastructure or components.

> **In scope — Platform Assets**
>
> - **Agent & wallet** — the Bankr agent runtime and its code sandbox, the Agent and Wallet APIs, and each account's gas-sponsored cross-chain wallet
> - **Token launches** — the launchpad: fair-launch token deployment, the swap-fee split, and creator-fee accounting
> - **Automations & orders** — conditional and scheduled trading the agent executes: limit and stop orders, DCA, TWAP, and recurring agent commands
> - **LLM Gateway** — the multi-provider LLM proxy at `llm.bankr.bot` : OpenAI- and Anthropic-compatible routing, per-key usage metering and credit billing, and tenant isolation
> - **x402 Cloud, webhooks & apps** — x402 paid-endpoint hosting and the payment flow, webhook delivery and signature verification, and the Apps platform of Bankr-generated mini-apps with server-side scripts
> - **Trading engine** — `swap.bankr.bot` and the engine that routes and executes external swaps and orders
> - **API keys & access controls** — API-key issuance and management, plus per-key controls like read-only mode and IP and recipient allowlists
> - **Bankr smart contracts** — deployed and owned by Bankr on **Base**:
>   **BankrOrders** — `0xE6A5762F31197BA43CE6C6E64709aC0679a95817`
>   **BankrFeeRouterV2** (x402) — `0x8AEE621035D93Deb3C0C1177fac252dC2dd501a0`

Out of scope — (*e.g.* not Bankr-operated or retired)

- **Third-party dependencies and integrations** — e.g. underlying DEX/router protocols, LLM providers behind the gateway, Privy, AWS, and other upstream providers. Report only as information or a courtesy to Bankr, and otherwise to the vendor; report here only if a Bankr-side misconfiguration or integration flaw creates the impact or is otherwise in scope. If unsure, contact security@bankr.bot.
- **User-launched tokens** — tokens and contracts deployed by users or their agents, unless a Bankr launchpad or platform flaw is the root cause. If unsure, contact security@bankr.bot.
- **User-side key & account compromise** — an API key or account compromised through the user's own device, phishing, malware, or accidental disclosure such as committing a key to a public repo is the user's responsibility, not a Bankr vulnerability. In scope only when a Bankr-side flaw exposes, leaks, or mishandles the key. If unsure, contact security@bankr.bot.
- **Lookalike / impersonation domains** — the only canonical domains are `bankr.bot`, `api.bankr.bot`, `docs.bankr.bot`, `help.bankr.bot`, `llm.bankr.bot`, `swap.bankr.bot`, and `x402.bankr.bot`; anything else is fake. Report clones so we can act, but they are abuse, not a platform vulnerability.
- **Infrastructure we do not own** and **social / third-party SaaS accounts**, unless a Bankr-side flaw is the root cause. If unsure, contact security@bankr.bot.
- **Installable / community skills** — third-party skill code is out of scope; report only a Bankr-side failure to isolate or validate it. If unsure, contact security@bankr.bot.
- **Staking v1 smart contract** — **BankrSecurityModule** `0x019fD9ABC9CaEB476D7AFA68BB675518C6BE17b7`, officially deprecated 8 April 2026; not in scope.

Bankr may modify the scope of this Policy, add or remove assets, or suspend the program at any time. Any changes apply prospectively; they do not affect rewards that Bankr has already approved in writing based on reports submitted before the change.

3 What Counts as a Vulnerability

For purposes of this Policy, a “vulnerability” is a weakness in Bankr-controlled systems or code that could allow an attacker to compromise the confidentiality, integrity, or availability of Bankr or user data or funds, even if no such compromise has yet occurred. A vulnerability lets an attacker do something they should not — access another user's data, move or freeze funds, escalate privileges, or degrade the platform for everyone. It must come with a reproducible exploitation path (§7).

A “cyber incident” is different: it involves actual or reasonably suspected loss or theft of confidential or personal information, unauthorized access to or use of systems or data, or malicious disruption of Bankr's systems. If your research leads you to access or obtain any personal information or other sensitive data, you must:

- Stop any testing that involves that data.
- Avoid saving, transferring, or using the data beyond what is strictly necessary to demonstrate the issue.
- Promptly report the vulnerability to Bankr and describe the type and amount of data you accessed, so we can assess any legal or contractual notification obligations.

Qualifying classes

- Authentication/authorization bypass, or unauthorized access to another user's account, wallet, orders, API keys, files, memory, or stored credentials (IDOR, broken access control, privilege escalation)
- Theft, freezing, or unauthorized movement of funds
- Prompt injection or agent-instruction manipulation that makes the agent move funds, leak sensitive data, or take unauthorized actions
- Bypass of the agent's transaction-security screening, or of wallet- and key-level controls (spend limits, recipient/IP allowlists, recipient cooldown, read-only mode, fail-closed), to move funds or exceed limits
- Leakage of secrets, keys, or credentials
- Remote code execution, SSRF, injection (SQL/command/template)
- Smart-contract or transaction-authorization flaws that drain, freeze, or misroute funds
- Tenant-isolation breaks in the LLM gateway exposing other tenants' keys or traffic
- Manipulation of LLM Gateway billing or credits — free or unmetered inference, credit theft, or bypassing the payment gate
- x402 payment-flow abuse — bypassing payment confirmation or settlement, defeating the allowedHosts allowlist (SSRF), or evading per-visitor rate limits
- Infrastructure or supply-chain compromise affecting the whole platform at once

Not rewarded on their own (without demonstrated impact)

- **DDoS / volumetric DoS** — not a vulnerability by itself (see §6)
- Best-practice nits with no demonstrated impact — missing security headers, weak TLS, outdated dependencies, and similar
- Automated scanner output with no working proof of concept
- Issues needing physical access, a fully compromised victim device, or social engineering of users or staff; self-XSS or clickjacking on non-sensitive pages
- Absence of rate limiting with no resulting account or fund impact
- Theoretical MEV/front-running or economic attacks with no concrete, quantified loss (or requiring capital irrational for the value at risk)

4 Vulnerability Categories & Severity

We group reports into two blast-radius tiers, then layer financial and data impact on top. Tier is the strongest single driver of reward.

Tier A — Single-user impact

Affects one user, or users one at a time through targeted exploitation. Examples:

- Compromise or leakage of an individual user's **API key** via a Bankr-side flaw
- IDOR letting an attacker read or modify one account's orders, alerts, or wallet metadata
- Permission flaw, session fixation, or token theft that lets one user act as or take over a single other account

Tier B — All-users / systemic impact

Affects every user at once, or the platform as a whole, without being repeated per victim — the highest-valued class. Examples:

- Compromise of a shared key, credential, or app secret
- Infrastructure or supply-chain compromise (CI/CD, deploy pipeline, private registry, exit node, gateway)
- A trading-engine or wallet flaw that drains or freezes funds across accounts
- Mass data or credential exposure — database-wide, or all tenants' keys and traffic via the LLM gateway

> **Rule of thumb:** if an attacker pays a fixed cost once and harms every user, it is Tier B; if the cost scales per victim, it is Tier A. A Tier A bug that *chains* into systemic impact is rated at the systemic level.

Severity

Final severity combines tier with financial and data impact, aligned with how crypto security researchers classify findings.

Severity	Typical profile	Representative impacts
Critical	Tier B with direct fund loss, or full systemic compromise	Direct theft of user/pooled funds; permanent fund freezing; protocol insolvency; key or credential compromise; RCE on core infrastructure
High	Tier B without direct loss, or Tier A with fund loss	Auth bypass affecting all users; theft of a single user's funds; temporary freezing of funds; platform-wide data exposure; authorization of attacker-chosen transactions in limited conditions
Medium	Tier A; sensitive data or account control, no direct loss	IDOR over one account's orders; single API-key leak; per-victim privilege escalation; griefing that degrades one user's operation
Low	Limited impact; requires unlikely preconditions	Minor information disclosure; low-value metadata leak with no account or fund impact

Raises severity	Lowers severity
Funds recoverable by attacker or unrecoverable by user; no user interaction; pre-auth (no valid account needed)	Requires a privileged or already-compromised account; needs significant victim interaction or social engineering

5 Reward Structure

Rewards under this program are discretionary, *ex gratia* payments that Bankr may choose to make to encourage good-faith security research and timely, responsible disclosure. Nothing in this Policy obligates Bankr to pay a reward for any particular report, even if it falls within a severity band described below. Bankr's decisions on eligibility, severity, and reward amount are final.

Rewards are paid on confirmed, in-scope, reproducible reports and scaled to demonstrated impact. Bankr sets the final amount; for Critical fund-loss reports it may scale with funds genuinely at risk. In setting the final amount, Bankr may consider factors such as the clarity and quality of the report, ease of exploitation, number and type of affected users, whether funds were genuinely at risk or actually lost, the complexity and cost of remediation, and whether the issue affects Bankr-specific code versus a broadly-shared component.

Severity	Reward band
Critical	$10k–$50k+
High	$2.5k–$10k
Medium	$500–$2.5k
Low	$100–$500

Reward conditions

- **First reporter.** Only the first researcher to report a previously unknown, reproducible issue is eligible. Reports that Bankr determines to be duplicates of an already-reported or already-known issue are not separately rewarded.
- **One reward per root cause.** Multiple reports arising from the same underlying flaw are treated as one; Bankr may combine related reports in its discretion.
- **Known and accepted-risk issues.** Reports of issues that Bankr has already identified and accepted as a risk (for example, documented limitations or trade-offs) generally do not qualify for rewards.
- **Caps and budgets.** Bankr may set caps on total rewards paid for a single underlying issue or vulnerability class, and on the total rewards payable to a single researcher or team, and may adjust program-level budgets over time.
- **Eligibility & compliance.** Payouts are subject to identity verification and compliance checks, including sanctions and AML screening. We cannot pay individuals or entities that are on applicable sanctions or restricted-party lists or located in comprehensively sanctioned jurisdictions, and we may require KYC before payout.
- **Anonymity.** You may submit reports anonymously or under a pseudonym. However, to be eligible for any monetary reward you must provide sufficient identity and contact information for Bankr to complete verification and compliance checks and to receive funds on Ethereum mainnet or Base.
- **Payment.** Rewards are paid in **USDC** on Ethereum mainnet or Base. After verification and screening, you'll provide a receiving address on one of those chains. Bankr is not responsible for loss of funds sent to an address you provide in error or that you do not control. You are solely responsible for any tax consequence or reporting obligations as a result of a payment. You are also responsible for cooperating with any Bankr request for information to complete required forms or documentation (e.g. 1099s) related to the same.

6 DDoS & Denial of Service

Denial of service is treated differently from other classes.

> **DDoS is not a vulnerability on its own**
>
> Volumetric or network-level denial of service — flooding endpoints, exhausting bandwidth, or overwhelming infrastructure — **does not qualify** on its own. It shows capacity, not a flaw in our code or logic.
>
> **The exception:** a DoS condition **does qualify** when it directly causes **loss of funds** — e.g. timed to block a liquidation, freeze in-flight orders, or strand balances. It is then valued on that financial impact, not the DoS itself.

Never run load tests or live DoS against production. Demonstrate the fund-loss mechanism logically or in an isolated or forked environment (not by attacking production), and describe the trigger precisely. Any DDoS event caused by you, whether or not in an attempt to qualify under this Policy, will not be rewarded and may result in corrective action or enforcement.

Network-level stress testing, capacity probing, traffic flooding, and similar activities are not authorized under this Policy. Do not intentionally generate traffic patterns designed to degrade service quality or exhaust resources.

7 Required: Clear Reproduction

Every report **must** include a clear, step-by-step reproduction. Reports without a working proof of concept may be closed as informational. Include:

1. **Summary** — one sentence: what the flaw is and its impact.
2. **Affected asset** — exact endpoint, domain, contract address, parameter, or component.
3. **Category & tier** — your assessment (Tier A single-user / Tier B all-users) and why.
4. **Preconditions** — what access or setup is needed (account, role, chain, balance).
5. **Steps to reproduce** — numbered, exact, copy-pasteable. Include requests, payloads, and parameters.
6. **Proof of concept** — script, curl, screenshots, or a short video. For contract issues, a runnable PoC on a fork or testnet — Anvil/Tenderly for EVM, devnet/localnet for Solana, testnet for Hyperliquid — never mainnet.
7. **Impact** — concretely what an attacker gains; quantify funds/users where possible.
8. **Suggested fix** — optional but appreciated.
9. **Contact information** — how to reach you and criteria necessary for payment of any reward.

8 Rules of Engagement

- **Do not** access, modify, or exfiltrate other users' data beyond the minimum needed to prove the issue.
- **Do not** run DoS/DDoS, spam, or high-volume automated scans against production.
- **Do not** move, drain, or test against real user funds — use your own test accounts and minimal amounts. Any attempt to move, freeze, or interfere with funds other than your own, beyond the minimum necessary to validate a vulnerability in a non-production or forked environment, is strictly prohibited and will be treated as abuse, not research, and is outside the scope of this Policy.
- **Do not** exploit contract issues on mainnet — demonstrate on a fork or testnet. No front-running or MEV against production users.
- **Do not** social-engineer Bankr staff, users, or vendors, or attempt physical intrusion. Social engineering of any kind (including phishing, pretexting, or baiting) against Bankr workforce members, users, or vendors is out of scope and voids safe harbor and any reward eligibility.
- **Do not** publicly disclose before we confirm a fix and agree on timing.
- **Do** stop and report as soon as you have confirmed impact.
- **Do** keep findings confidential until coordinated disclosure (§9).

Good-faith research within these rules will not be pursued as a violation (see §10). Acting outside scope — data theft, fund movement, extortion, or disruption — voids safe harbor and any reward and may lead to separate action by Bankr.

9 Disclosure Process & Response Targets

Stage	What happens	Target
Submit	Email security@bankr.bot or file a report at Bankr Help Center (Request Type: **Bug Bounty / Vulnerability**), with the full report (§7).	—
Acknowledge	We confirm receipt of your report.	2 business days
Triage	We validate reproduction and assign category + severity.	5–10 business days
Remediate	We fix and may ask you to re-test; timeline scales with severity.	Severity-dependent
Reward	Paid on confirmed, in-scope issues after verification & screening (§5).	Post-fix
Disclose	Coordinated public disclosure after the fix ships, by mutual agreement.	By agreement

These timelines are targets, not guarantees. Complex issues, dependencies on third-party vendors, multi-party coordination, or integration with incident-response and regulatory processes may require longer periods. Bankr will make good-faith efforts to keep you informed of material delays.

In some cases, Bankr may be legally or contractually required to notify regulators, partners, or affected users about certain vulnerabilities or related incidents. Bankr may make such notifications independently of any coordinated public disclosure we agree with you.

We aim to keep researchers updated at each stage. If a report seems stalled, reply to the thread for an update.

10 Authorization, Safe Harbor & Legal Action

This section defines the boundary between authorized research and unlawful conduct. The protections below apply **only** to activity within this policy.

Good Faith Security Research

For purposes of this Policy, “good faith security research” means accessing or testing in-scope Bankr systems solely to investigate, identify, and help correct security vulnerabilities, in a manner designed to avoid harm to Bankr, our users, and the public, and using any information derived from testing only to:

- promote the security of Bankr's systems and users; or
- inform Bankr or relevant developers of the vulnerability so it can be fixed.

Research is not in good faith if it involves extortion or ransom demands, sale or transfer of exploit details or exfiltrated data to third parties, exploitation beyond what is necessary to demonstrate a vulnerability, or any activity outside the Rules of Engagement (§8), scope (§2), or coordinated disclosure process (§9).

Where research is conducted in good faith and full compliance with this policy — including the Rules of Engagement (§8) and timely, confidential disclosure to security@bankr.bot or Bankr Help Center (§9) — Bankr treats it as authorized and will not pursue or support legal action for it.

Authorization and Safe Harbor

When you perform good faith security research strictly within this Policy — including the Scope (§2), What Counts as a Vulnerability (§3), Required Reproduction (§7), Rules of Engagement (§8), and Disclosure Process (§9) — Bankr will treat your testing as authorized for purposes of computer-misuse and anti-circumvention laws, to the fullest extent permitted by law. In particular, for such in-scope, compliant research, Bankr:

- Considers your access and testing to be “authorized” under the U.S. Computer Fraud and Abuse Act (CFAA) and similar state computer-misuse laws, to the extent those laws apply.
- Does not intend to pursue civil action or initiate a law-enforcement complaint for accidental, good-faith violations of this Policy that occur in the course of such research.
- If a third party alleges that you violated computer-misuse laws in connection with research that we can verify complied with this Policy, will make known, where reasonably appropriate, that your actions were conducted under Bankr's vulnerability disclosure program and within its stated limitations.

Similarly, for good faith security research conducted within this Policy, Bankr does not intend to bring or support claims against you under the anti-circumvention provisions of the Digital Millennium Copyright Act (DMCA), provided that your testing is performed in an environment designed to avoid harm to individuals or the public and information derived from your testing is used solely to improve security or inform Bankr.

This safe harbor does not extend to activity on systems or assets that are out of scope, to any research that involves actual theft, destruction, or widespread disruption, or to any actions that violate other applicable laws. We cannot, and do not, authorize any activity that would violate laws or regulations in your applicable jurisdiction.

> **Abuse, Non-Disclosure, and Extortion**
>
> Conduct outside this policy is not authorized and forfeits safe harbor and any reward eligibility. Bankr may pursue civil or criminal remedies, or cooperate with law enforcement or regulators, against anyone who:
>
> - intentionally abuses, exploits, or attacks Bankr's systems or users;
> - discovers a vulnerability and weaponizes it or fails to disclose it promptly and confidentially in accordance with this Policy;
> - attempts to move or misappropriate funds or data; or
> - engages in extortion, ransom demands, or any threat to disclose, sell, or weaponize a vulnerability or exfiltrated data in exchange for payment or other advantage.
>
> Bankr will not negotiate or pay ransom or extortion demands, and will not re-label ransom payments as “bug bounties.” Bankr also reserves all rights and remedies under applicable law, including computer-misuse, fraud, data-protection, theft, and sanctions statutes in any relevant jurisdiction. Nothing in this Policy waives any rights except the limited safe harbor for good faith security research described above.
>
> Anyone who intentionally abuses, exploits, or attacks Bankr's systems — or who discovers a vulnerability and fails to disclose it under this policy — will be subject to legal action, and **Bankr will pursue all available civil and criminal remedies to the fullest extent of the law.**

Without limitation, the following forfeit any reward and safe-harbor protection:

- Any violation of the Rules of Engagement (§8) — including accessing another user's data, moving funds, or disclosing publicly before coordinated disclosure (§9);
- Discovering a vulnerability and failing to disclose it promptly and confidentially, or exploiting it for any purpose other than a good-faith proof of concept;
- Extortion, ransom demands, or any threat to disclose or weaponize a finding;
- Selling, transferring, or sharing a vulnerability or exfiltrated data with any third party, or any access or action exceeding the authorization granted by this policy.

Program Changes and Relationship to Other Terms

Bankr may update or discontinue this Policy, adjust scope, severity models, or rewards, and terminate any individual's participation in the program at any time, in each case on a prospective basis. Bankr will not terminate a participant's eligibility solely to avoid paying a reward that Bankr has already agreed in writing to pay for a qualifying report.

From time to time, Bankr may operate private or invite-only programs that require participating researchers to accept additional written terms (for example, non-disclosure agreements or participation terms). Where such terms apply, they will control to the extent of any conflict with this public Policy.

11 Miscellaneous

This Policy does not create any enforceable rights or obligations and does not itself form a contract between Bankr and any researcher. To the extent any legal relationship is deemed to arise from participation in this program, it will be governed by the laws of the State of Delaware, excluding its conflict-of-laws principles. If any provision of this Policy is held invalid or unenforceable, the remaining provisions will remain in full force and effect.