JOROMI AI

ACCEPTABLE USE POLICY

Rules Governing the Responsible, Lawful, and Secure Use of the JOROMI AI Enterprise Operating System

Version [●] — Effective as of 14th July

This policy is incorporated by reference into the JOROMI AI Terms of Service, Enterprise Subscription Agreement, and Master Subscription Agreement / Enterprise Services Agreement, and should be read alongside the Privacy Policy, Data Processing Agreement, Responsible AI Policy, and Information Security Policy.

1. Introduction

This Acceptable Use Policy (this "Policy") sets out the rules that govern the use of the JOROMI AI Enterprise Operating System, including its AI Agents, Multi-Agent Systems, AI Agent Studio, Workflow Automation, Enterprise Search, Knowledge Management, Document Intelligence, AI Copilots, Business Intelligence, Analytics, Governance and Compliance tooling, APIs, Identity and Access Management, Cloud Infrastructure, Retrieval-Augmented Generation capabilities, Enterprise Knowledge Bases, and AI Automation Studio (collectively, the "Platform", and the services delivered through it, the "Services"), made available by Joromi Technologies Inc. or its applicable contracting affiliate ("JOROMI AI," "we," "us," or "our").

Purpose. This Policy exists to protect the integrity, security, and reliability of the Platform; to protect the rights and safety of Authorized Users, Data Subjects, and third parties; to promote responsible and ethical use of artificial intelligence; and to maintain the trust that enterprise, government, and regulated customers place in JOROMI AI when they deploy the Platform across sensitive business functions. A shared and clearly understood set of rules allows organizations of every size — from a five-person startup to a multinational financial institution — to build on the Platform with confidence.

Responsible AI Principles. JOROMI AI is built around the premise that AI Systems are most valuable, and least risky, when they are deployed with meaningful human oversight, appropriate transparency, and respect for the rights of the people whose data and decisions they touch. This Policy translates those principles into specific, enforceable rules, and works alongside JOROMI AI's Responsible AI Policy, which describes JOROMI AI's own design commitments in greater depth.

Platform Integrity and Enterprise Trust. Because many organizations rely on the Platform concurrently, the misuse of the Platform by one organization or individual can affect the security, performance, or reputation of the Platform for others. This Policy is designed to prevent that kind of harm, and JOROMI AI enforces it consistently across all customers, regardless of size or contract value.

This consistency matters most for the customers who rely on it least visibly — a government agency conducting sensitive casework, a hospital system processing patient records, or a bank running compliance workflows all depend on the same underlying platform integrity as a small business automating its invoicing. This Policy is written with that full range of customers in mind, and its enforcement does not vary based on the sophistication or negotiating leverage of the customer involved.

Cybersecurity and Ethical Technology Use. This Policy reflects JOROMI AI's commitment to cybersecurity best practices and to the ethical use of technology more broadly, including prohibitions on activity that would facilitate cyberattacks, fraud, discrimination, or harm to vulnerable individuals, as detailed in Sections 5 and 6.

Relationship with Other JOROMI AI Documents. This Policy is incorporated by reference into, and forms part of, the Terms of Service, the Enterprise Subscription Agreement, and any Master Subscription Agreement or Enterprise Services Agreement between JOROMI AI and a Customer (collectively, the "Principal Agreement"). It should be read together with the Privacy Policy, the Data Processing Agreement ("DPA"), the Responsible AI Policy, and the Information Security Policy, each of which addresses related subject matter in greater detail. In the event of a direct conflict between this Policy and the Principal Agreement regarding permitted or prohibited use of the Platform, this Policy governs solely with respect to that subject matter, and the order of precedence set out in the Principal Agreement otherwise applies.

How to Read This Policy. This Policy is organized so that each audience can find what applies to it without reading the entire document from end to end: Section 4 describes what Customers and Authorized Users may generally do with the Platform; Sections 5 and 6 describe what is prohibited, with Section 6 focused specifically on the particular risks that arise from generative and agentic AI functionality; Sections 7 through 9 describe the affirmative responsibilities that Customers and Authorized Users take on when they use the Platform; and Sections 10 through 12 describe how JOROMI AI monitors for, investigates, and responds to violations, and how anyone — whether a Customer, an Authorized User, or an unaffiliated third party — can report a concern. Nothing about this structure changes the binding effect of any individual provision; every section of this Policy is enforceable according to its terms.

2. Scope and Applicability

This Policy applies to every individual and organization that accesses or uses the Platform in any capacity, including:

- Customers that have entered into a Principal Agreement with JOROMI AI;
- Authorized Users designated by a Customer to access the Services under that Customer's Account;
- Enterprise Administrators responsible for configuring Workspaces, provisioning users, and managing security settings;
- Developers and API Users who integrate with the Platform programmatically;
- contractors, vendors, and consultants engaged by a Customer who are granted access to the Platform;
- partners and third-party integrators that connect their own applications or services to the Platform;
- Marketplace developers who publish AI Agents, connectors, or extensions for use by other Customers;
- educational institutions and their students, faculty, and staff granted access under an Education plan;
- government agencies and their personnel granted access under a Government plan;
- trial users evaluating the Platform prior to entering into a paid subscription; and
- participants in any beta or early access program made available by JOROMI AI.

Organizational Responsibility. Each Customer is responsible for ensuring that its Authorized Users, employees, contractors, and any other individual to whom it grants access to the Platform comply with this Policy. A Customer's failure to ensure such compliance may result in enforcement action against the Customer's Account in accordance with Section 11, regardless of whether the Customer itself directly engaged in the violation.

Marketplace Developers and Third-Party Integrators. Where JOROMI AI operates a marketplace or directory through which developers publish AI Agents, connectors, templates, or other extensions for use by other Customers, this Policy applies to the published extension itself, and the publishing developer is responsible for ensuring that the extension does not facilitate, encourage, or automate any activity prohibited under Section 5 or Section 6. JOROMI AI may remove a published extension, suspend a developer's publishing privileges, or take other enforcement action under Section 11 if an extension violates this Policy, independent of any action taken against the underlying developer's own Account.

Trial and Beta Access. Individuals and organizations accessing the Platform on a trial basis, or as participants in a beta or early access program, are subject to this Policy in the same manner as paying Customers,

notwithstanding the more limited scope of functionality, support, or service levels that may apply to such access under the Principal Agreement or a separate beta program agreement.

Geographic Scope. This Policy applies to use of the Platform regardless of the jurisdiction from which it occurs. Where a specific provision of this Policy references a particular law or regulatory framework, that provision applies in addition to, and does not narrow, the general prohibitions and responsibilities set out elsewhere in this Policy for Customers and Authorized Users located in other jurisdictions.

3. Definitions

Capitalized terms used in this Policy have the meanings set out below, or, where not defined here, the meanings given in the Principal Agreement.

"Account" means the registration record created for a Customer, or for an individual Authorized User, to access the Services.

"AI Agent" means a configurable, semi-autonomous software entity built on an AI System that performs tasks, makes retrieval calls, or executes workflow steps on behalf of a Customer with limited or no direct human intervention at each step.

"AI Output" means any content, prediction, classification, summary, recommendation, code, or other material generated by an AI System in response to an AI Prompt or as part of an automated workflow or AI Agent task.

"AI Prompt" means any instruction, query, document, or other input submitted to an AI System, whether directly by an Authorized User or programmatically, to obtain AI Output.

"API" means an application programming interface made available by JOROMI AI for programmatic access to the Services.

"Authorized User" means an individual whom a Customer has authorized to access and use the Services under that Customer's Account.

"Customer" means the entity that has entered into a Principal Agreement with JOROMI AI for access to the Services.

"Customer Data" means all data, content, files, and information submitted to, or generated within, the Services by or on behalf of a Customer.

"Enterprise Workspace" means a configurable environment within a Customer's Account in which Authorized Users organize Customer Data, AI Agents, and workflows.

"Platform" has the meaning given in Section 1.

"Services" has the meaning given in Section 1.

"Security Incident" means any confirmed unauthorized or unlawful access to, acquisition of, use of, disclosure of, alteration of, or destruction of Customer Data or Platform infrastructure.

"Malicious Code" means any virus, worm, trojan horse, ransomware, spyware, logic bomb, or other code, script, or program designed to disrupt, damage, or gain unauthorized access to a computer system or data.

"Sensitive Data" means Personal Data revealing racial or ethnic origin, health information, biometric or genetic data used for identification purposes, financial account credentials, government-issued identification

numbers, precise geolocation data, or other categories of data that Applicable Data Protection Laws treat as warranting heightened protection.

"Prohibited Content" means content that violates Section 5 or Section 6 of this Policy, including unlawful, infringing, fraudulent, or harmful content as further described therein.

"Abuse" means any use of the Platform that violates this Policy, degrades the security, performance, or availability of the Platform, or is intended to circumvent technical or contractual limitations applicable to a Customer's subscription.

"Reverse Engineering" means decompiling, disassembling, or otherwise attempting to derive the source code, underlying model weights, training methodology, or architecture of the Software or any AI System.

"Intellectual Property Rights" means all patent, copyright, trademark, trade secret, moral rights, and other intellectual property rights recognized in any jurisdiction worldwide, whether registered or unregistered.

"Data Subject" means an identified or identifiable natural person to whom Personal Data contained within Customer Data relates.

"Beta Feature" means any feature of the Platform designated by JOROMI AI as "beta," "preview," "early access," or by similar terminology, indicating that the feature is made available for evaluation purposes and has not reached general availability.

"Trial User" means an individual or organization accessing the Services on a trial basis prior to entering into a paid subscription under an Order Form.

"Marketplace" means any directory, catalog, or storefront operated by JOROMI AI through which developers may publish AI Agents, connectors, templates, or other extensions for use by Customers.

4. Permitted Use

This Section 4 is intended to give Customers and Authorized Users a practical sense of the breadth of legitimate activity the Platform is designed to support, so that the prohibitions in Sections 5 and 6 are read in proper context: they are boundaries around a very wide range of permitted, encouraged, and commercially valuable use, not a narrow carve-out from a generally restrictive policy.

Subject to the terms of the Principal Agreement and this Policy, Customers and Authorized Users may use the Platform for lawful, legitimate business, governmental, educational, and research purposes, including:

- AI-assisted business operations, including drafting, summarization, and task automation performed with appropriate human oversight;
- enterprise productivity and internal collaboration among Authorized Users within a Workspace;
- knowledge management, including the creation and maintenance of Enterprise Knowledge Bases;
- workflow automation of routine or repetitive internal business processes;
- research and analysis conducted in accordance with applicable ethical and legal standards;
- analytics and business intelligence derived from Customer Data;
- software development using the Platform's APIs, developer tools, and AI Agent Studio;
- document intelligence, including extraction, classification, and summarization of enterprise documents;

- enterprise governance and internal compliance monitoring;
- supporting regulatory compliance efforts, including audit trail generation and policy enforcement;
- educational instruction and coursework at educational institutions granted access under an Education plan;
- innovation and prototyping of new AI-enabled internal tools and workflows;
- customer service and support operations directed at an organization's own customers; and
- AI experimentation conducted within the lawful and contractually authorized limits of a Customer's subscription, including in sandbox or non-production environments where made available.

All use of the Platform, however permitted in principle under this Section 4, must at all times comply with applicable law, the Principal Agreement, this Policy, and any internal policies the Customer has adopted governing its own personnel's use of AI tools.

Sandbox and Non-Production Use. Where JOROMI AI makes sandbox, developer, or non-production environments available as part of a Customer's subscription, Customers may use those environments to experiment with AI Agent configurations, test integrations, and evaluate new workflows before deploying them into production, provided that such experimentation remains within the Usage Limits and any environment-specific restrictions communicated by JOROMI AI, and does not involve the submission of production Customer Data or Sensitive Data unless the sandbox environment is specifically designed and approved for that purpose.

5. General Prohibited Activities

This Section 5 sets out categories of conduct that are prohibited on the Platform regardless of whether artificial intelligence functionality is involved. Section 6 supplements this Section 5 with additional restrictions specific to the use of AI Systems, AI Agents, and AI Output. The categories below are illustrative rather than exhaustive: conduct that is substantially similar in nature or effect to an enumerated prohibition, even if not specifically listed, is also prohibited.

Customers and Authorized Users must not use the Platform to engage in, facilitate, or promote any of the following:

5.1 Unlawful and Harmful Conduct

- violating any applicable law, regulation, or third-party right in any jurisdiction in which the Platform is used;
- infringing or misappropriating any Intellectual Property Right of JOROMI AI or any third party;
- committing or facilitating fraud, identity theft, or misrepresentation of identity or authority;
- money laundering, terrorist financing, or other financial crime;
- human trafficking, forced labor, or exploitation of any individual;
- the sexual exploitation or abuse of minors in any form, which is strictly prohibited under all circumstances and will be reported to the appropriate authorities;
- harassment, stalking, intimidation, or threats directed at any individual or group;

- hate speech or content that promotes discrimination, violence, or hostility against individuals or groups based on a protected characteristic;
- defamatory content that a Customer knows or reasonably should know to be false; and
- the creation, storage, or distribution of content that is otherwise illegal under applicable law, including unlawful weapons, controlled substances, or child sexual abuse material.

5.2 Technical Abuse and Security Violations

- developing, uploading, or distributing Malicious Code through the Platform;
- conducting or facilitating ransomware attacks, phishing campaigns, or credential theft schemes;
- sending unsolicited bulk communications (spam) through or using the Platform;
- gaining or attempting to gain unauthorized access to any Account, system, or data not belonging to the accessing party;
- bypassing, disabling, or attempting to circumvent any security control, authentication mechanism, or technical limitation of the Platform;
- disrupting or degrading the network, infrastructure, or performance of the Platform, including through denial-of-service attacks or excessive automated requests;
- engaging in Reverse Engineering of the Software or any AI System, except to the extent such restriction is prohibited by applicable law;
- conducting unauthorized benchmarking, penetration testing, or vulnerability scanning of the Platform without JOROMI AI's prior written authorization;
- circumventing subscription tier limitations, Usage Limits, or API rate limits through technical workarounds, credential sharing, or automated evasion techniques; and
- scraping, crawling, or systematically harvesting data from the Platform outside of supported APIs and integrations.

5.3 Commercial Misuse

- misrepresenting a Customer's identity, size, or affiliation in connection with its use of the Platform;
- reselling, sublicensing, or otherwise commercially exploiting access to the Platform beyond the scope of the license granted under the Principal Agreement; and
- using the Platform to build, train, or benchmark a directly competing product or service, or to conduct competitive analysis for that purpose, except as separately agreed in writing with JOROMI AI.

These commercial-use restrictions are intended to protect the substantial investment JOROMI AI makes in the Platform and its underlying AI Systems, and do not restrict a Customer's ordinary use of AI Output that it lawfully generates for its own internal business purposes or for delivery to its own customers as part of a product or service that is not itself substantially similar to the Platform.

6. AI-Specific Restrictions

Generative and agentic AI functionality creates value precisely because it can produce fluent, convincing, and often highly capable output from limited instruction — and that same capability can be misdirected toward deception, manipulation, or automation of harmful conduct at a scale that would be difficult to achieve manually. This Section 6 addresses those AI-specific risks directly, in addition to the general prohibitions in Section 5, which continue to apply in full to any use of the Platform's AI functionality.

In addition to the general prohibitions in Section 5, and given the particular risks associated with generative and agentic AI functionality, Customers and Authorized Users must not use AI Systems, AI Agents, AI Prompts, or AI Output to:

- generate content that is itself unlawful, including fraudulent financial or legal documents, counterfeit credentials, or fabricated identification documents;
- create fake identities, sock-puppet accounts, or synthetic personas intended to deceive individuals or platforms;
- produce deepfakes, synthetic media, or voice clones of real individuals intended to deceive, defraud, or cause reputational or financial harm, without the consent of the individual depicted;
- conduct social engineering, pretexting, or other manipulation intended to induce an individual to disclose credentials, Sensitive Data, or confidential information;
- plan, coordinate, or execute cyberattacks, including through AI-generated malware, exploit code, or attack infrastructure;
- develop or refine exploits targeting known or unknown vulnerabilities in software, hardware, or networks;
- conduct unauthorized surveillance of individuals, including through covert location tracking, communications interception, or biometric identification without a lawful basis and appropriate consent;
- misuse biometric data, including facial recognition or voiceprint data, in a manner inconsistent with Applicable Data Protection Laws or without an appropriate legal basis;
- build manipulative AI systems designed to exploit cognitive biases, deceive users about the nature of their interaction with AI, or covertly influence individuals' decisions against their own interests;
- interfere with or seek to manipulate the outcome of an election or other democratic process, including through disinformation generated or amplified using AI Systems;
- engage in market manipulation, including the AI-assisted generation or dissemination of false or misleading information intended to influence securities or commodities prices;
- make high-risk autonomous decisions — including decisions materially affecting an individual's employment, credit, insurance, housing, or legal rights — without meaningful human review appropriate to the risk, consistent with Section 6.1;
- extract, replicate, or attempt to reconstruct the underlying weights, architecture, or training methodology of any AI System made available through the Platform;
- conduct prompt injection attacks against AI Agents or AI Systems, whether the Customer's own or those of another Customer, or attempt to manipulate an AI System into disregarding its configured instructions or safety controls; or

- attempt to "jailbreak" an AI System or otherwise circumvent content or safety controls built into the Platform.

6.1 Human Review of High-Impact AI Output

Customers remain responsible for reviewing AI Output before relying on it in connection with legal, financial, medical, regulatory, employment, or other high-impact decisions affecting any individual. AI Systems may generate inaccurate, incomplete, or biased content, including content that appears authoritative but is factually incorrect. JOROMI AI's Responsible AI Policy and the AI-specific provisions of the DPA and Enterprise Subscription Agreement describe this requirement in further detail, and this Policy should be read consistently with those documents.

6.2 Autonomy Limits for AI Agents

Where a Customer configures an AI Agent to take actions with real-world or system-level effect — such as sending communications on the Customer's behalf, modifying records in a connected system, executing financial transactions, or triggering downstream automated workflows — the Customer is responsible for configuring approval checkpoints, permission scopes, and monitoring appropriate to the potential consequences of the AI Agent's actions. Customers must not configure an AI Agent to take irreversible, high-consequence actions entirely without human checkpoints where a reasonable organization would insist on such oversight, such as the wholesale deletion of production data, the execution of material financial transactions, or communications directed at regulators or the public that purport to bind the Customer.

6.3 Marketplace AI Agents and Third-Party Models

AI Agents, connectors, or model integrations published by third-party developers or accessed through the Platform's marketplace are subject to the same restrictions set out in this Section 6, and a Customer's use of a third-party AI Agent does not relieve the Customer of its own obligations under this Policy, including the human review obligations described in Section 6.1.

6.4 Reporting Suspected Misuse of AI Functionality

Given the pace at which techniques for misusing generative AI evolve, JOROMI AI treats reports of suspected AI-specific misuse — including newly identified jailbreak techniques, prompt injection vectors, or emerging patterns of AI-facilitated fraud — as a priority input into how this Section 6 and the Platform's underlying safety controls are updated over time. Customers and Authorized Users who identify such a technique, whether through their own testing or through observing misuse by others, are encouraged to report it through the channels described in Section 12 rather than publicly disclosing it in a manner that could facilitate its use by bad actors before JOROMI AI has an opportunity to respond.

7. Data and Privacy Responsibilities

JOROMI AI acts as a Processor of Customer Personal Data submitted to the Platform, and the Customer, as Controller, retains primary responsibility for the lawfulness of that data's collection and use, as further described in the DPA. This Section 7 summarizes the practical steps a Customer should take to meet that responsibility in the context of everyday Platform use.

Customers are responsible for ensuring that their use of the Platform complies with all applicable privacy and data protection laws, including with respect to:

- establishing a lawful basis for collecting and submitting Personal Data to the Platform;
- obtaining any consents required under applicable law before submitting Personal Data, including Sensitive Data, to the Platform;
- applying data minimization principles, submitting only the Personal Data reasonably necessary for the Customer's intended purpose;
- maintaining the confidentiality of Customer Data, including by configuring appropriate access permissions within Enterprise Workspaces;
- implementing the security controls described in Section 8 to protect Customer Data submitted to the Platform;
- maintaining the accuracy and quality of Customer Data submitted to the Platform;
- granting Platform permissions and access rights consistent with the principle of least privilege described in Section 8;
- complying with applicable privacy laws in any cross-border transfer of Personal Data facilitated through the Platform; and
- handling Sensitive Data submitted to the Platform with heightened care, including limiting its collection to circumstances where it is genuinely necessary and legally permitted.

Further detail regarding the respective data protection obligations of JOROMI AI and its Customers is set out in the DPA, which governs the Processing of Personal Data by JOROMI AI as Processor on behalf of the Customer.

Sensitive Data Warning for AI Features. Customers should exercise particular caution before submitting Sensitive Data as an AI Prompt or as part of Enterprise Data indexed for Retrieval-Augmented Generation, since such data may subsequently surface in AI Output to any Authorized User with access to the relevant Workspace, Enterprise Knowledge Base, or AI Agent. Customers should configure Workspace permissions accordingly and should consult their own legal or compliance function before submitting categories of Sensitive Data that carry heightened regulatory risk, such as health records or biometric identifiers.

8. Security Responsibilities

Security on the Platform is a shared responsibility. JOROMI AI is responsible for the security of the Platform itself, as described in the Information Security Policy and the DPA; Customers and Authorized Users are responsible for the security practices described in this Section 8, which govern how they configure and interact with the Platform.

Customers and Authorized Users are responsible for maintaining appropriate security practices in connection with their use of the Platform, including:

- using strong, unique passwords for Platform Accounts and not sharing credentials among individuals;
- enabling multi-factor authentication (MFA) where made available, particularly for Administrator and other privileged Accounts;

- maintaining accurate identity management records and promptly de-provisioning Accounts for personnel who no longer require access;
- configuring access controls consistent with the principle of least privilege, granting Authorized Users only the permissions necessary to perform their roles;
- regularly reviewing user access and permissions within Enterprise Workspaces to identify and remove unnecessary or excessive access;
- maintaining reasonable device security, including up-to-date operating systems and endpoint protection, on devices used to access the Platform;
- promptly reporting any suspected Security Incident to JOROMI AI in accordance with Section 12;
- responsibly disclosing any suspected vulnerability in the Platform to JOROMI AI's Security Team rather than exploiting or publicly disclosing it prior to remediation; and
- configuring integrations and API credentials securely, including rotating API keys periodically and restricting their use to authorized systems.

9. Customer Responsibilities

Beyond the specific data, privacy, and security obligations set out in Sections 7 and 8, Customers take on a broader set of organizational responsibilities when they deploy the Platform across their business. These responsibilities reflect the reality that JOROMI AI provides the underlying capability, but the Customer determines how that capability is actually used within its own organization.

In addition to the data, privacy, and security responsibilities described above, each Customer is responsible for:

- managing its Authorized Users, including timely provisioning and de-provisioning of Accounts;
- establishing internal governance processes appropriate to its use of AI Agents, AI Output, and automated workflows;
- training its employees and other Authorized Users on the appropriate and permitted use of the Platform, including this Policy;
- maintaining reasonable backup practices for Customer Data that the Customer does not wish to rely on the Platform alone to preserve;
- monitoring its own compliance with applicable law and internal policy in its use of the Platform;
- reviewing AI-generated outputs and applying appropriate human oversight before relying on them for consequential decisions, as described in Section 6.1;
- promptly reporting any known or suspected misuse of the Platform by its own Authorized Users or third parties, in accordance with Section 12; and
- cooperating with JOROMI AI's reasonable requests for information in connection with an investigation into suspected violations of this Policy.

A Customer that operates in a regulated industry — such as financial services, healthcare, or government — should assess whether its use of the Platform requires additional internal controls, disclosures, or approvals beyond those described in this Policy, and should not treat this Policy as a substitute for its own sector-specific compliance obligations.

10. Platform Monitoring

JOROMI AI may monitor activity on the Platform, including through automated systems and, where necessary, manual review, for the purposes of detecting abuse, preventing fraud, investigating suspected Security Incidents, protecting the integrity and availability of the Platform, enforcing the Principal Agreement and this Policy, and complying with applicable legal obligations.

Monitoring conducted under this Section 10 is carried out in a manner consistent with the Privacy Policy and applicable law, and is not intended to, and does not, grant JOROMI AI a general right to review the substantive content of Customer Data absent a legitimate basis described in this Section 10 or the DPA. Where monitoring reveals evidence of a suspected violation of this Policy, JOROMI AI may take enforcement action in accordance with Section 11.

JOROMI AI's monitoring relies primarily on automated systems that analyze patterns of activity — such as anomalous API request volumes, known indicators of Malicious Code, or signatures associated with previously identified abuse — rather than routine manual review of Customer Data. Manual review of specific content is generally reserved for circumstances in which automated systems, a report received under Section 12, or another legitimate basis gives JOROMI AI reason to investigate further.

11. Enforcement

If JOROMI AI becomes aware of a suspected violation of this Policy, it may take any of the following actions, individually or in combination, as appropriate to the nature and severity of the suspected violation:

- conduct an investigation, including review of relevant logs, Account activity, and, where necessary and permitted, Customer Data directly relevant to the suspected violation;
- impose temporary restrictions on the affected Account or specific Platform functionality pending completion of an investigation;
- suspend the affected Account, in accordance with the suspension provisions of the Principal Agreement;
- suspend the affected Services generally, where necessary to protect the Platform or other Customers;
- terminate the affected Account or Principal Agreement, in accordance with the termination provisions of the Principal Agreement, for severe, repeated, or uncured violations;
- remove or disable access to specific content, AI Agents, or workflows implicated in the violation;
- restrict or revoke API access associated with the violation;
- report the conduct to relevant regulators or governmental authorities where required or appropriate, particularly in connection with suspected child exploitation, terrorism, or other serious crimes;
- report the conduct to law enforcement where JOROMI AI reasonably believes a crime has been or is being committed;
- preserve evidence relevant to the suspected violation, including for the purpose of cooperating with a law enforcement or regulatory investigation; and
- take any emergency measures reasonably necessary to prevent imminent harm to the Platform, other Customers, or third parties, including immediate suspension without prior notice.

Where appropriate given the nature of the suspected violation, JOROMI AI will seek to take enforcement action proportionate to the severity of the violation, and, where consistent with protecting the Platform and its Customers, will provide notice and an opportunity to remediate before imposing the most severe enforcement measures. JOROMI AI is not obligated to provide advance notice where doing so would risk harm to the Platform, other Customers, or third parties, or where required by law.

11.1 Appeals

A Customer that disputes an enforcement action taken under this Section 11 may request reconsideration by contacting the team identified in Section 12 that is most relevant to the nature of the action (for example, the Abuse Team for a content-related suspension, or the Security Team for an action taken in response to a suspected Security Incident), providing any information relevant to the dispute. JOROMI AI will review such requests in good faith, but reconsideration of an enforcement action does not suspend its effect unless JOROMI AI expressly agrees otherwise in writing.

11.2 Data Handling Following Enforcement Action

Where an Account is suspended or terminated under this Section 11, the handling of the affected Customer's Customer Data following that suspension or termination is governed by the data retention and deletion provisions of the DPA and the Enterprise Subscription Agreement, except that JOROMI AI may retain data reasonably necessary to investigate the violation, to comply with a legal obligation, or to support a law enforcement or regulatory referral made under this Section 11, for so long as reasonably necessary for those purposes.

12. Reporting Violations

JOROMI AI encourages Customers, Authorized Users, and third parties to report suspected violations of this Policy, including:

- suspected Security Incidents or vulnerabilities affecting the Platform;
- suspected abuse or misuse of the Platform by any Customer or Authorized User;
- suspected intellectual property infringement involving content hosted on or generated through the Platform;
- suspected fraud connected to use of the Platform;
- suspected unauthorized access to any Account or Customer Data; and
- any other suspected violation of this Policy.

Reports may be submitted through the following channels:

- Abuse Team: [Insert abuse report email address / web portal URL]
- Security Team / Security Operations Center: [Insert security incident reporting email address, and, if applicable, PGP key or secure reporting portal for vulnerability disclosure]
- Legal Department: [Insert legal department contact email address]
- Ethics Hotline: [Insert ethics hotline phone number and/or web portal URL, including whether anonymous reporting is supported]

- General Web Portal: [Insert URL of general trust and safety reporting portal]

JOROMI AI will review reports submitted in good faith and take appropriate action consistent with Section 11. JOROMI AI will not retaliate against any individual for making a good-faith report of a suspected violation of this Policy, and will treat the identity of a reporting individual as confidential to the extent consistent with conducting a thorough investigation and complying with applicable law.

JOROMI AI aims to acknowledge receipt of reports involving suspected Security Incidents or imminent harm within a reasonably prompt timeframe appropriate to the severity of the report, recognizing that a report alleging an active security threat or ongoing harm to an individual warrants faster escalation than a routine policy question. Reports that do not include sufficient detail to investigate may result in a request for additional information before JOROMI AI can take further action.

13. Compliance Framework

This Policy is designed to support alignment with major privacy, security, and AI governance frameworks, including the GDPR, the UK GDPR, the Nigeria Data Protection Act (NDPA), the CCPA/CPRA, the NIST Cybersecurity Framework, applicable AI governance frameworks, and, where applicable to JOROMI AI's certification status, ISO/IEC 27001, ISO/IEC 27701, and SOC 2 principles.

Nothing in this Section 13 constitutes a representation that JOROMI AI holds any specific certification, has received any specific regulatory approval, or is in full compliance with any specific framework, except to the extent expressly confirmed in JOROMI AI's then-current Security Documentation or a written statement signed by an authorized representative of JOROMI AI.

Customers operating in jurisdictions with sector-specific or emerging AI-specific regulation — including, for example, financial services supervisory requirements, healthcare data protection rules, or AI risk-management obligations applicable to a particular industry — should independently assess whether this Policy and JOROMI AI's other documentation are sufficient to satisfy those requirements, and should engage their own legal counsel where appropriate. JOROMI AI's Compliance Team, identified in Section 15, can assist with such assessments by providing available documentation upon request.

14. Policy Updates

JOROMI AI reviews this Policy periodically and may update it to reflect changes in applicable law, industry practice, emerging risks associated with AI technology, or the functionality of the Platform. Each version of this Policy is identified by a version number and effective date at the top of the document.

For changes that materially reduce Authorized Users' rights or materially expand the scope of prohibited conduct, JOROMI AI will provide notice through the methods it customarily uses to communicate with Customers, which may include email notice to Administrators, in-product notification, or posting of the updated Policy with a summary of material changes. Non-material clarifications and updates may be made without advance notice.

Continued use of the Platform following the effective date of an updated version of this Policy constitutes acceptance of that updated version. Customers that do not agree to a material update may raise concerns through their designated account contact or, where applicable, exercise any termination rights available under the Principal Agreement.

15. Contact Information

Questions regarding this Policy, or reports of suspected violations, may be directed to:

- Abuse Team: [Insert email address]
- Security Operations Center (SOC): [Insert email address / phone number]
- Privacy Office: [Insert email address]
- Legal Department: [Insert email address]
- Compliance Team: [Insert email address]
- Ethics Office: [Insert email address / hotline number]
- Website: [Insert corporate website URL]
- Registered Office Address: [Insert registered office address]