

Lythe Acceptable Use Policy

Last updated: 21 October 2026

1. Purpose

This Acceptable Use Policy ("Policy") sets out the rules for accessing or using Lythe products and services provided by Lythe Inc. or its applicable affiliate, including Lythe Pte. Ltd. (collectively and as applicable, "Lythe"), including Fraser Simulation Studio, websites, applications, APIs, dashboards, public demos, AI personas, subagents, simulations, automated evaluation workflows, agent orchestration tools, visual evaluation systems, bug-testing workflows, integrations, documentation, and related services (collectively, the "Services").

At launch, Lythe Inc. provides the Services to business customers in the United States. Lythe Pte. Ltd. is the contracting provider only where expressly identified in a separately executed customer agreement. Public access to informational websites does not establish eligibility for the Services or restrict rights under applicable privacy law.

The purpose of this Policy is to prevent misuse of the Services, protect End Users, customers, and third parties, support lawful business research and product evaluation, reduce fraud and spam, prevent misleading synthetic research, and maintain safe, reliable, and trusted AI-powered simulation and evaluation products.

Capitalized terms used but not defined in this Policy have the meanings given to them in the Lythe AI Use Policy, applicable Terms of Service, applicable Terms of Use, or other applicable agreement.

2. Scope

This Policy applies to all customers, users, administrators, developers, contractors, agents, affiliates, and any person or entity that accesses or uses the Services.

Customers are responsible for the conduct of their account users, representatives and third parties they authorize to use their account, and for their requests, instructions, permitted materials, customer-controlled systems and integrations, and downstream use of Outputs. Lythe is responsible for configuring, executing, monitoring and stopping its own agents and simulations within accepted Run Records and applicable safeguards.

This Policy supplements the applicable Lythe Terms of Service, Terms of Use, Lythe AI Use Policy, Lythe Cookie Policy, Lythe Privacy Policy, Order Forms, customer agreements, and any other applicable product-specific terms that expressly incorporate or reference this Policy.

Mandatory transfer terms prevail over conflicting provisions for the relevant transfer. The DPA prevails over conflicting provisions concerning Customer Personal Data Processing. Subject to those rules, an executed Master Services Agreement, Order Form, or other signed customer

agreement prevails for its expressly stated subject matter, followed by the Terms of Service, the Acceptable Use Policy and AI Use Policy, and other incorporated policies. Target Certifications and feature authorizations narrow permitted use and cannot override mandatory law, required safeguards, or the controlling agreement. For eligible non-enterprise purchases, the Refund Policy has the specific priority provided in Section 24 of the Terms of Service.

Nothing in this Policy limits or modifies the obligations of Lythe Inc. or Lythe Pte. Ltd., as applicable, under any applicable Privacy Policy, Data Processing Addendum, data processing agreement, or applicable data protection law.

3. Customer Responsibilities

Customers must ensure their requests, genuine testing authority, Run Confirmations, permitted inputs, systems they control and downstream use comply with applicable law, industry rules, platform terms and internal policies. Customers specify their needs, objectives, requested targets and scenarios through the vendor or client portal or another enabled request interface. Lythe manually verifies testing authority, configures and runs accepted simulations within the agreed plan and guardrails, monitors and stops its agents, and delivers agreed reports. A request does not directly execute a simulation; required Customer confirmations remain conditions of execution.

Use the Services only for lawful, authorized, and legitimate business purposes.

Obtain and maintain required permissions, notices, consents, lawful bases and source records before collecting and supplying Permitted Inputs, connecting customer-controlled integrations or requesting simulations. Supply genuine testing authority and fresh Run Confirmations before execution. Real customer records, feedback or conversation materials require Data Approval before receipt or access; target authority, payment and general feature approval do not replace it.

Ensure that AI personas, subagents, simulations, and outputs do not misrepresent who or what they are, who they represent, the source of feedback, or the purpose of an interaction or evaluation.

Use accurate, current and lawfully obtained source materials, and supply only Permitted Inputs as defined in the Terms of Service. Public-page personal details must be ordinary, incidental and necessary for the approved evaluation. Real customer records, feedback or conversation materials require Data Approval before receipt or access; Lythe recommends synthetic sample records for authenticated testing. Safeguarded test accounts remain required; real customer records must remain hidden unless Lythe expressly approves the particular records, target access and safeguards in a Data Approval before receiving or accessing them.

Identify relevant review requirements and provide an authorized contact for questions, escalations and stop requests. Lythe must configure appropriate operational human escalation for the sensitive, high-risk, uncertain or failed simulations and workflows it controls; Customers must provide appropriate review for downstream uses they control.

Review delivered reports and Outputs before downstream reliance and promptly report known complaints, anomalies, policy violations, unreliable or harmful Outputs and misleading synthetic research. Lythe must monitor its agents, simulations and AI personas and investigate operational anomalies and material concerns.

Maintain security controls for account access, API keys, customer-controlled integrations and source systems, Permitted Inputs, uploaded materials, generated reports, and any unexpected personal data pending isolation, return, or deletion.

Promptly notify Lythe of suspected misuse, unauthorized access, security incidents, privacy issues, legal complaints, or harmful behavior related to the Services.

4. Prohibited Uses

Customers and users may not use the Services, or allow the Services to be used, for any of the following activities.

4.1 Illegal, Harmful, or Abusive Activity

Any activity that violates applicable laws, regulations, court orders, sanctions, export controls, industry rules, platform policies, carrier rules, or third-party rights.

Harassment, intimidation, threats, stalking, coercion, hate, abuse, exploitation, targeted humiliation, or non-consensual surveillance.

Fraud, scams, phishing, impersonation, identity theft, extortion, blackmail, or deceptive schemes.

Attempts to obtain sensitive information through deception, including passwords, one-time codes, account credentials, banking details, government identifiers, authentication secrets, or private security information.

Activities that facilitate violence, self-harm, terrorism, human trafficking, child exploitation, or other serious harm.

Generating or distributing malware, unauthorized code, credential theft tools, or instructions for compromising systems.

Calls, messages, simulations, outputs, reports, or automated workflows that are unlawful, unsolicited, misleading, coercive, deceptive, abusive, or harmful.

4.2 Spam, Unwanted Communications, and Contact Data Misuse

Sending or initiating spam, unsolicited emails, messages, or other unwanted communications using the Services or outputs.

Failing to comply with applicable data-protection, privacy, and communication laws regarding research communications, including required notices, consents, and lawful bases for collecting or using personal data.

Contacting individuals or data subjects without required consent, notice, lawful basis, business relationship, or other authorization required by applicable law.

Using purchased, scraped, stolen, rented, misleading, outdated, or unlawfully obtained contact lists or research data.

Ignoring opt-outs, revocation of consent, suppression lists, deletion requests, or unsubscribe requests.

Using the Services, outputs, or reports to evade compliance rules, regulatory requirements, carrier rules, or platform enforcement.

Using misleading identities, spoofing names, rotating domain names, or concealing the origin or sponsor of research or evaluations.

Conducting high-volume unsolicited outreach or research communications without appropriate compliance review, lawful bases, consent records, and opt-out mechanisms.

Contacting individuals on do-not-contact registries, suppression lists, or other opt-out mechanisms in any applicable jurisdiction.

Using the Services for sales outreach, marketing communications, promotional campaigns, fundraising, or outbound lead generation unless expressly approved in writing by Lythe and conducted in compliance with applicable law, platform requirements, and Lythe policies.

Disguising sales, marketing, lead generation, fundraising, promotional, or debt-collection communications as feedback, support, technical support, customer success, research, or requested evaluation communications.

Soliciting, collecting, or processing personal data through deceptive, coercive, or misleading means.

4.3 Misrepresentation, Impersonation, and Deceptive AI

Making an AI persona, subagent, simulation, or automated evaluator pretend to be a specific real person without authorization.

Making an AI persona, subagent, simulation, or automated evaluator falsely claim to be human where disclosure is required or where non-disclosure would be deceptive or harmful.

Impersonating government officials, emergency services, banks, healthcare providers, legal representatives, employers, public agencies, customers, experts, public figures, protected groups, or other trusted institutions without authorization.

Misrepresenting the customer, sponsor, identity, research methodology, source of feedback, purpose of the simulation or evaluation, product, pricing, legal effect, or recipient obligations.

Using names, likenesses, identities, or customer personas without required consent and rights.

Using Lythe services to manipulate, pressure, or deceive vulnerable individuals, including older adults, minors, patients, consumers in financial distress, or people in emergency situations.

Creating, presenting, or distributing simulated feedback, AI persona responses, outputs, evaluation results, or reports with false human provenance or concealed synthetic origin, including as feedback from real users or real customer interviews; presenting outputs as guaranteed market truth; or making claims of independently validated research, statistically valid survey results, or validated market findings that do not meet Section 6.2.

4.4 Regulated and High-Risk Activities

The Services may not be used for regulated or high-risk activities without Lythe's prior written approval, unless Lythe expressly states in writing that prior approval is not required for a particular category of use. Any approved deployment must be supported by appropriate controls, contracts, disclosures, compliance documentation, human oversight, and legal review appropriate to the proposed use.

Medical diagnosis, emergency medical triage, clinical decision-making, or providing medical treatment instructions without qualified human oversight.

Legal advice, immigration advice, tax advice, financial advice, investment recommendations, lending decisions, insurance underwriting, or eligibility determinations without qualified human oversight and required licenses.

Credit, housing, employment, education, insurance, healthcare, public benefits, criminal justice, or other consequential decisions about individuals without appropriate legal basis, human review, auditability, and appeal mechanisms.

Political campaigning, voter persuasion, fundraising, polling, or election-related calls, messages, simulations, or targeting where disclosure, consent, registration, or caller-identification rules apply, unless approved in writing and configured to comply with applicable law.

Debt collection, repossession, collections outreach, or enforcement activity without written approval and compliance controls.

Emergency services, crisis hotlines, law enforcement dispatch, or safety-critical systems where failure, delay, hallucination, or incorrect output could cause serious harm.

4.5 Sensitive Data Misuse

Collecting, requesting, inferring, storing, uploading or using sensitive or special-category personal data, children's data, real payment-card credentials, re-identification keys or other information prohibited by the Terms of Service; or using ordinary personal information outside Permitted Inputs.

Requesting real-person passwords, authentication codes, payment-card details, government identifiers, biometric identifiers, precise individual location, health information, or other sensitive data through tests. Approved test-account credentials may be handled only through the secure mechanism permitted by the Terms of Service. Lythe enables authenticated testing only after Customer provides a test environment or test accounts with safeguards appropriate to the target.

Uploading, supplying or exposing real customer records, personal feedback, recordings, transcripts or conversation materials before a specific Data Approval, or outside its approved scope. Identifiable raw replay remains prohibited even if approval is requested. Necessary incidental ordinary personal details on authorized public pages are subject to the separate public-page limits. Do not upload material Customer has no right or lawful basis to process.

Using call transcripts, recordings, CRM records, customer data, research data, simulation outputs, or AI persona responses to discriminate, profile, exploit, or target individuals unlawfully.

Attempting to identify or re-identify individuals, reverse anonymization, supply re-identification keys, or combine inputs or Outputs to create profiles of identifiable individuals.

4.6 Platform Abuse, Security Abuse, and Reverse Engineering

Interfering with, disrupting, overloading, scanning, probing, scraping, crawling, attacking, or compromising the Services or related infrastructure.

Bypassing rate limits, usage caps, authentication, access controls, monitoring systems, safety filters, billing controls, or compliance controls. The restrictions in this Policy on bypassing an Authorized Target's security, authentication or access controls are subject only to the narrow written-authorization exception in Sections 6.5 and 7 of the Terms of Service. All other restrictions continue to apply.

Sharing, selling, exposing, or failing to secure API keys, access tokens, credentials, secrets, call data, simulation data, uploaded materials, or integration credentials.

Reverse engineering, copying, decompiling, or attempting to extract source code, model prompts, system prompts, agent orchestration logic, simulation methods, evaluation methods, scoring methods, or proprietary systems except where law expressly permits.

Using the Services, outputs, evaluation results, prompts, workflows, or proprietary materials to benchmark, train, improve, or build a competing product in violation of an applicable agreement or Lythe's intellectual property or contractual rights.

Using automated systems to create accounts, simulate usage, generate artificial simulation traffic, or manipulate billing, analytics, reputation, deliverability, conversion metrics, or platform metrics.

Using public demos or Services to crawl, scrape, test, evaluate, overload, or analyze third-party websites, applications, prototypes, products, or systems without authorization.

4.7 Content and Intellectual Property Misuse

Uploading, generating, transmitting, storing, or analyzing content that infringes intellectual property, privacy, publicity, contractual, confidentiality, or other rights.

Using copyrighted materials, proprietary scripts, recordings, voices, datasets, screenshots, websites, product materials, user research materials, or customer records without required rights.

Removing ownership notices, watermarking, audit markers, compliance disclosures, or provenance information.

Using Lythe services to create defamatory, deceptive, illegal, infringing, or harmful content.

5. Call and Conversation Data Requirements

Customers may supply synthetic or truly anonymized datasets for labeling and persona enrichment. Real customer feedback, conversation-derived material or transcripts may be supplied only after Lythe separately approves a Data Approval before receipt or access. Recordings additionally require an enabled feature and specific approval. Approved material may be used for that Customer's own synthetic personas through prompts, retrieval or configuration, without model training, profiling real people or use of identifiable material for other customers. Remove unnecessary identifiers and prohibited sensitive material before transfer.

5.1 Consent and Notices for Call and Conversation Data

Customers must have all required rights, notices, consents, lawful bases and source authorizations for collecting and supplying permitted feedback and conversation materials. Data Approval must cover any ordinary personal information retained; anonymous inputs must be genuinely anonymized, including relevant indirect identifiers.

Customers must comply with applicable call recording, privacy, and data-protection laws (including one-party, two-party, all-party, or other consent rules) in relevant jurisdictions.

Customers must maintain records of consent, lawful bases, call purposes, and required disclosures where required.

Identifiable feedback, transcripts, conversation material or call recordings must not reach Lythe before a specific Data Approval. Recordings require an enabled feature and approval covering the recording format and necessary ordinary personal information. Required recording notices or consent do not permit sensitive information or replace the other Data Approval requirements. Voice recognition or biometric identification is not authorized.

5.2 Disclosure and Transparency for Persona Enrichment

Customers must accurately disclose the purpose of call and conversation data usage and must not misrepresent or conceal that call or conversation data may be processed by AI personas, simulations, or automated evaluators.

Customers must not request AI personas, subagents, simulations or automated evaluators that deny they are AI, falsely claim to be human or mislead individuals about automation. Lythe must not configure its agents to engage in that conduct.

Customers must use clear, accurate, and non-deceptive language about persona identity, sponsor, purpose, capabilities, and data usage.

Customers must provide a clear path to appropriate human review or escalation for their downstream uses where required by law, contract or the nature of that use. Lythe must provide appropriate human review and escalation for the simulation execution it controls.

Customers must not request, or present Outputs in a way that conceals, the automated nature of AI personas, subagents, simulations or evaluators where a reasonable person would be misled, deceived or materially influenced. Lythe must not configure its agents to facilitate that concealment.

5.3 Human Oversight and Safety

Lythe must provide qualified human review or appropriate escalation for sensitive, uncertain, disputed, high-risk or failed simulations and evaluations it operates. Customers must identify relevant requirements, respond to necessary questions and apply appropriate human review to downstream decisions they control.

Customers must not use AI personas, subagents, simulations, or automated evaluators to make final determinations in regulated, safety-critical, or consequential contexts unless expressly approved in writing by Lythe and supported by applicable legal authority, qualified human oversight, appropriate controls, and required documentation.

Customers must review delivered Outputs before reliance, report known inaccurate, unsafe, discriminatory or non-compliant results to Lythe, and correct affected downstream uses they

control. Lythe must investigate reported concerns and correct or stop affected persona or evaluator behavior within its control as appropriate.

Customers must not rely on the Services as the sole system for emergencies, crisis response, or legally required human communications.

The Services are not designed for emergency calling, crisis response, law enforcement dispatch, medical emergencies, or any situation where delay, failure, or incorrect output could result in death, injury, or serious harm.

6. Fraser Simulation Studio Specific Requirements

Because Fraser Simulation Studio enables AI personas, simulated product feedback, automated evaluations, and product research workflows, customers must follow additional requirements.

6.1 Authorized Targets, Inputs, and Permissions

Customers may request Fraser Simulation Studio testing only of an Authorized Target as defined in the Terms of Service. Before Lythe accesses or tests a target, Customer must submit a Target Certification identifying the exact target and scope. If Customer does not own the target, Customer must possess prior express written authorization from the Target Owner specifically permitting Customer and Lythe to perform the applicable automated, AI-agent or simulation testing. Lythe's manual authority review and each required fresh Customer Run Confirmation must be completed before execution.

Customers must not request testing of a competitor or other third-party target based only on public availability, a hyperlink, a merchant, vendor, customer or partner relationship, ordinary user access, credentials, an API key, a robots.txt setting or an absence of technical blocking. Customers must retain authorization evidence, provide it as required before execution, remain within its scope and immediately notify Lythe and withdraw affected instructions through the designated stop route if authority expires, is withdrawn, becomes disputed or is exceeded. Lythe must stop affected work under the accepted stop procedure.

Customers must obtain required rights, notices, consents, lawful bases and source permissions for all inputs. Only Permitted Inputs may reach Lythe through URLs, screenshots, feedback, conversation material, datasets, APKs or integrations. Real customer records, feedback and conversation materials require Data Approval before receipt or access. Lythe recommends synthetic sample records for authenticated testing. Safeguarded test accounts remain required; real customer records must remain hidden unless Lythe expressly approves the particular records, target access and safeguards in a Data Approval before receiving or accessing them. Confidential or executable materials require appropriate authorization and safeguards. Sensitive or otherwise prohibited information and malicious or unlawful materials remain prohibited.

Testing must remain low-volume and non-destructive within the certified and approved scope. Authenticated testing and customer-specific integrations require the feature and data permissions in the Terms of Service. Lythe enables authenticated testing only in client-provided test environments or with safeguarded client-provided test accounts appropriate to the target. Lythe recommends synthetic sample records for authenticated testing. Safeguarded test accounts remain required; real customer records must remain hidden unless Lythe expressly approves the particular records, target access and safeguards in a Data Approval before receiving or accessing them. At launch, Customers must not request or use tests to contact real people or perform real purchases, payments or other real transactions, and Lythe must not execute those actions; written approval of a target or data category does not override those restrictions. Customer must maintain safeguards for test accounts and environments it controls and immediately report a failure or scope breach through the designated stop route. Lythe must enforce the accepted scope and stop affected execution if safeguards fail or scope is exceeded. Other restricted activities, including security testing, vulnerability scanning, production-data changes, load testing or destructive activity, remain subject to the Terms of Service. Target authorization and data permission are separate requirements, and changes outside the approved scope require renewed certification or approval before affected activity begins.

6.2 Synthetic Feedback and Research Claims

Customers must clearly disclose the synthetic origin of AI personas, simulated users, automated evaluations, synthetic feedback, simulation reports, and outputs. Simulated participants, interviews, and feedback must never be represented as real people, real customer interviews, or feedback from real users, even if substantive findings are independently validated. Claims of independently validated research, statistically valid survey results, or validated market findings are permitted only to the extent supported by appropriate independent research methods, with the methodology and synthetic contribution disclosed. Outputs must not be presented as guaranteed market truth.

Customers must use clear, accurate, and non-deceptive language when describing simulation outputs, assumptions, limitations, research methodology, evaluation criteria, and the synthetic nature of AI persona feedback.

Customers must not use AI personas or simulations to impersonate a specific real person, customer, employee, expert, public figure, or protected group without authorization.

6.3 Reliance and Human Review

Customers must not rely solely on simulation outputs for regulated, high-impact, consequential, safety-critical, product-launch, investment, legal, compliance, employment, housing, insurance, financial, medical, or legally binding decisions.

Customers must independently verify product critiques, simulated feedback, classifications, recommendations, reports, and other outputs before relying on them for consequential, regulated, or legally significant actions.

Customers must identify requested evaluation goals, success criteria, prohibited outcomes and relevant downstream constraints. Lythe must define and enforce failure conditions, operational review, escalation and stopping or rollback procedures for simulations and systems it controls. Customers remain responsible for review and rollback procedures for changes or downstream uses they control. All testing remains subject to accepted Run Records, required Run Confirmations and launch restrictions.

7. Prohibited Content, Prompts, and Configurations

Customers may not submit prompts, scripts, instructions, datasets, knowledge base content, call flows, voices, examples, simulations, evaluation criteria, AI persona configurations, or integrations that are designed to cause Lythe services to violate this Policy, bypass safety controls, deceive users, misrepresent synthetic feedback, or produce prohibited outputs.

This includes prompt injection, jailbreak attempts, hidden instructions, adversarial inputs, or tool misuse intended to evade Lythe's own access controls, safety requirements, billing controls, monitoring, or approved testing scope; unauthorized persona creation; misleading research instructions; unauthorized website testing; and false or misleading Target Certifications. Evaluation of an Authorized Target using adversarial test inputs remains subject to the Target Owner's express authorization, Lythe's required written approval for security testing, and all restrictions in the Terms of Service. This paragraph does not grant an independent testing permission.

8. Customer Data, Contact Lists, and Uploaded Materials

Customers represent and warrant that they have required rights, notices, consents, authorizations and lawful bases to collect source materials and supply Permitted Inputs. Necessary public-page personal details are limited to approved evaluation, and real customer records, feedback or conversation materials require Data Approval before receipt or access. Identifiable source datasets or authenticated-target records outside that prior approval, and sensitive or otherwise prohibited inputs, must not be supplied. Lythe generally recommends synthetic sample records instead of real customer records.

Do not upload stolen, scraped, purchased, unlawfully obtained, or unauthorized customer data, research data, call recordings, transcripts, product materials, datasets, APK files, or third-party content.

Data subject to recording, confidentiality or other special restrictions requires the applicable legal authority, notices, consents and safeguards and must remain within Permitted Inputs and

any Data Approval. Sensitive or regulated personal data and children's data remain prohibited as testing or persona-enrichment inputs. Review or approval cannot waive those prohibitions.

Promptly delete or suppress contacts who opt out, revoke consent, or request deletion where required.

Keep CRM, integration, workspace, and product permissions limited to what is necessary for the intended workflow.

Do not use the Services to identify, re-identify, profile, or target identifiable individuals from customer inputs or Outputs.

Do not use simulation outputs or AI persona responses to make unsupported claims about real users, market demand, product-market fit, customer intent, investment outcomes, or customer behavior.

9. Monitoring, Enforcement, and Remediation

Lythe may monitor use of the Services to detect abuse, enforce this Policy, protect the platform, comply with law, and respond to complaints. Monitoring may include review of account metadata, usage patterns, reports, complaints, logs, prompts, scripts, simulation settings, evaluation criteria, transcripts, recordings, conversation metadata, uploaded materials, outputs, and reports, subject to applicable law and customer agreements.

If Lythe believes that a customer or user has violated this Policy, Lythe may take enforcement action with or without prior notice where necessary to prevent harm, comply with law, protect the Services, respond to complaints, or prevent abuse.

Enforcement actions may include:

Requesting additional information, compliance evidence, consent records, proof of rights, data-source details, workflow documentation, or research methodology documentation.

Requiring changes to prompts, disclosures, workflows, integrations, AI personas, subagents, simulation settings, evaluation criteria, uploaded materials, or usage practices.

Suspending or limiting specific APIs, features, integrations, simulations, reports, workspaces, uploads, or accounts.

Blocking traffic, prompts, files, uploads, simulations, or outputs.

Removing, quarantining, or disabling content, data, configurations, reports, or workflows.

Reporting suspected illegal activity to appropriate authorities, service providers, affected parties, or third-party platforms where required or appropriate.

Terminating access to the Services.

Seeking indemnification, damages, or other remedies available under contract or law.

Lythe may require proof of consent, data-source details, rights documentation, workflow descriptions, simulation purpose, business identity information, recording notices, AI disclosures, human-review workflows, research methodology, evaluation criteria, and other compliance documentation before enabling or continuing any workflow, simulation, upload, or use case.

Customer indemnification obligations are set out in the applicable Terms of Service, Master Services Agreement, Order Form, or other written agreement with the applicable Lythe entity.

10. Compliance Review and High-Volume or High-Risk Use

Lythe may require compliance review before enabling or continuing certain use cases, including regulated industry workflows, political or public-sector evaluations, healthcare, financial services, insurance, debt collection, government services, processing or analyzing large volumes of conversation-derived information, high-volume simulations or evaluations, public-facing AI persona outputs, third-party website, application, APK, or digital asset testing, externally relied-on research outputs, or unusual traffic patterns. Compliance review does not itself replace a specific Data Approval, authorize sensitive or otherwise prohibited inputs, or authorize a target; required Target Owner authority remains necessary.

Customers must cooperate with reasonable compliance reviews and provide accurate information about their use case, lawful basis, consents or notices for call audio or conversation data processing, data processing requirements, simulation purpose, uploaded materials, research methodology, evaluation criteria, and human-review or escalation controls.

Customers must not launch, continue, or materially change any use case requiring compliance review until Lythe has provided written approval.

Lythe does not provide legal, regulatory, research, statistical, or compliance advice. Customers are responsible for obtaining their own legal, regulatory, research, and compliance advice regarding their use of the Services.

11. Reporting Abuse

Anyone may report suspected abuse, illegal activity, spam, impersonation, unauthorized call recordings or conversation data, misleading synthetic research, unauthorized website or digital asset testing, privacy concerns, security issues, or violations of this Policy to Lythe using the contacts below.

Abuse contact: team@lythe.ai

Security contact: team@lythe.ai

Privacy contact: team@lythe.ai

Reports of concerns should identify relevant accounts, targets, request references and the issue without copying personal feedback, raw recordings, transcripts, credentials, prohibited data or re-identification keys into ordinary tickets or email. Where evidence containing permitted personal information is necessary, use the agreed secure route and access restrictions.

The Lythe entity responsible for a customer's use of the Services is identified in the applicable Terms of Service, Order Form, customer agreement, or other written agreement.

12. Changes to This Policy

Lythe may update this Policy from time to time to reflect changes in the Services, legal developments, telephony requirements, AI requirements, carrier or platform requirements, safety learnings, risk controls, business practices, research abuse patterns, or other abuse patterns. Updated versions will be posted or otherwise made available to affected customers and users.

Material changes to this Policy that affect an existing Customer's contractual rights or obligations are subject to the notice, affirmative acceptance and existing-purchase protections in Section 23 of the Terms of Service or the controlling executed customer agreement. Silence, account activity or continued use alone does not constitute the required acceptance. Emergency suspension or restrictions already permitted by the governing agreement may take effect for safety or legal reasons without expanding testing permission. Minor corrections that do not alter rights or obligations may be notified without renewed acceptance. Any separate consent required by applicable law must be obtained before the relevant Processing begins.