

Lythe AI Use Policy

Last updated: 21 October 2026

1. Purpose

This AI Use Policy explains the rules that apply when customers, developers, End Users, contractors, and other users access or use Lythe products and services provided by Lythe Inc. or its applicable affiliate, including Lythe Pte. Ltd. (collectively and as applicable, “Lythe”). This includes Fraser Simulation Studio, AI personas, subagents, simulations, automated evaluation tools, bug-testing workflows, visual evaluations, Simulation Evaluation tools, APIs, integrations, dashboards, transcripts, analytics, public demos, and related services (collectively, the “Services”). Only Permitted Inputs as defined in the Terms of Service may be supplied or exposed to Lythe. These include synthetic or truly anonymized information, necessary ordinary personal details incidental to authorized public-page testing, and real customer records, feedback or conversation materials separately approved through a Data Approval completed 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. Approved feedback may enrich only that Customer’s own synthetic personas through prompts, retrieval or configuration, without model training or profiling real people. The Services do not provide general-purpose outbound telephony, voice agents, or marketing campaigns. Service notifications, customer-selected integration updates, and email within approved test workflows, where enabled, remain subject to the Terms of Service’s target, feature, and communication restrictions. This description does not authorize unsolicited outreach or real-user communications. At launch, the prohibition on tests contacting real people or performing real transactions applies regardless of feature or target approval.

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 help ensure that the Services are used safely, lawfully, transparently, and reliably. Users may not use the Services to deceive people; misrepresent synthetic feedback as real user research; violate applicable privacy, data protection, consumer protection, intellectual property, or AI disclosure laws; deceptively impersonate any person, business, government entity, institution, customer, expert, or protected group; manipulate vulnerable individuals; or deploy AI personas, subagents, simulations, or automated evaluators without appropriate oversight, disclosure, and controls; or test, crawl, scrape, or monitor third-party websites, applications, systems, prototypes, or APK files without authorization.

2. Scope

This policy applies to all access to and use of the Services, whether through the web application, APIs, SDKs, integrations, simulation tools, developer portals, dashboards, public demos, internal tools, beta features, customer pilots, enterprise trials, or production deployments.

Customers are responsible for their account users, requested objectives, targets and scenarios, genuine testing authority, Run Confirmations, permitted inputs and instructions, customer-controlled systems and integrations, and downstream use of Outputs. Customers must ensure those matters comply with applicable laws, regulations, contractual obligations and this policy. Lythe configures and operates its own AI personas, subagents, simulations and automated evaluators within accepted Run Records and is responsible for the execution it controls. Customers do not operate Lythe's agents by submitting a request.

Developers are responsible for the products, workflows and integrations they control, including appropriate disclosures, permissions, safeguards, monitoring and human oversight for those systems and their downstream use of Outputs. This does not transfer responsibility for operating Lythe's agents or simulation systems to Customer.

These responsibilities apply in addition to, and do not limit, Lythe's own obligations under applicable law.

Lythe may investigate, suspend, restrict, or terminate access to the Services where use violates this policy or creates legal, safety, security, compliance, reputational, or platform-integrity risk.

3. Definitions

For purposes of this AI Use Policy, the following terms have the meanings set out below:

“AI Agent” means an automated or semi-automated system that receives inputs, reasons over context, generates outputs, takes actions, or interacts with individuals or systems through text, APIs, software tools, or other interfaces.

“AI Persona” means an AI-generated or AI-assisted simulated user, customer profile, evaluator, reviewer, tester, respondent or behavioral model that Lythe configures and operates using permitted Customer requirements and inputs to test, review, critique or analyze Customer Data, products, websites, user flows, prototypes, positioning, messaging, customer journeys or related materials.

“Simulation Evaluation” means the use of multiple agents, AI personas, subagents, simulations, tests, red-teams, automated evaluators, reviewers, or similar evaluation systems to test, compare, monitor, improve, or stress-test the behavior, performance, reliability, or safety of AI personas, subagents, simulations, automated evaluators, product experiences, user flows, messaging, positioning, customer journeys, or related workflows.

“Customer Data,” “Permitted Inputs” and “Data Approval” have the meanings in the Terms of Service. Customer Data includes permitted prompts, materials, public-page content, screenshots, feedback or conversation information, configurations, logs and Outputs processed on Customer’s behalf. Personal information within Permitted Inputs is protected by the DPA; information outside those boundaries is handled only for containment, investigation, secure return or deletion.

“End User” means any person who interacts with an AI persona, simulation, public demo, or other Lythe service, including leads, customers, patients, constituents, applicants, employees, research subjects, users, or other individuals.

“Output” means responses, transcripts, classifications, recommendations, analytics, evaluations, scores, simulated feedback, product critiques, insights, reports, actions, or other content generated by the Services, AI personas, or simulations.

“High-Impact Decision” means any decision, recommendation, assessment, ranking, determination, or action that may materially affect an individual’s employment, credit, housing, healthcare, insurance, education, legal status, financial access, public benefits, safety, or access to similar rights, services, or opportunities.

“Regulated Use Case” means any use of the Services involving legal, medical, financial, insurance, housing, employment, education, government, political, children or minors, emergency, or other activities or contexts subject to heightened legal, regulatory, professional, or safety requirements.

4. Core Principles

4.1 Transparency: Customers must ensure that End Users understand when they are interacting with an AI system where disclosure is required by applicable law, regulation, platform policy, customer agreement, Lythe policy, or where non-disclosure would be materially deceptive.

4.2 Lawful Sources: Customers must obtain and maintain required rights, permissions, notices, consents and lawful bases for collecting source materials and supplying Permitted Inputs. Real customer records, feedback or conversation materials require prior Data Approval. Customer ownership, target permission, consent or feature approval does not independently authorize data outside the permitted boundaries.

4.3 Human Accountability: Customers remain responsible for their requests, authority, permitted inputs, instructions and downstream reliance on Outputs. Lythe remains responsible for configuring and operating its AI personas, subagents, simulations and automated evaluators within accepted Run Records and applicable safeguards.

4.4 Safety by Design: Lythe must appropriately scope, test, monitor, log and supervise the AI personas, subagents, simulations and automated evaluators it operates and provide appropriate

human escalation. Customers must supply accurate requirements and valid authority, maintain safeguards for systems they control and promptly report concerns or request stopping through the designated route. Required Customer Run Confirmations remain conditions of execution.

4.5 No Deception: The Services must not be used to mislead people about identity, authority, capability, intent, sponsorship, research methodology, source of feedback, or material facts.

4.6 Data Minimization: Customers and Lythe must limit information to what is necessary for the approved workflow. Use synthetic or truly anonymized inputs where feasible, remove unnecessary identifiers and sensitive material before transfer, and minimize screenshots, logs and model inputs. Real customer records, feedback and conversation materials must remain within their prior Data Approval. Labeling is not permission to accept unapproved source data or to identify people; re-identification keys are prohibited.

5. Permitted Uses

Every listed use is conditional on the feature being available, the Terms of Service permitting the activity and data, and any required written approval. This list does not override restrictions on personal data, production testing, transactions, or communications. Subject to those conditions, legitimate business workflows may include:

Labeling synthetic or truly anonymized datasets and analyzing conversation-derived information or feedback within a prior Data Approval to configure or enrich that Customer's own synthetic personas and simulations through prompts, retrieval or configuration, without model training or profiling real individuals.

Internal business operations such as research preparation, QA reviews, workflow automation, and CRM updating where proper consent exists.

Customer support, frequently asked questions, ticket creation, and status updates.

Business analytics, process improvement, and other internal operations.

Product feedback simulations, landing page reviews, user flow testing, onboarding critiques, positioning analysis, messaging analysis, feature evaluation, product-market-fit hypotheses, bug testing, and visual evaluations of websites, applications, and product flows.

Agent testing, simulation, red-teaming, prompt evaluation, performance improvement, and reliability testing through Simulation Evaluation.

AI simulation pilots in regulated industries, provided that customers implement appropriate compliance controls, obtain any required internal approvals, maintain appropriate human oversight, and obtain any prior written approval required from Lythe under Section 14.

6. Prohibited Uses

Users may not use Lythe services, including Fraser Simulation Studio, for any of the following activities:

6.1 Illegal, Harmful, or Abusive Activity

Violating any applicable law, regulation, court order, industry rule, platform policy, or contractual obligation.

Fraud, scams, phishing, social engineering, identity theft, credential theft, account takeover, or deceptive collection of personal data.

Harassment, threats, intimidation, extortion, hate, abuse, stalking, doxxing, or non-consensual surveillance.

Generating or facilitating malware, unauthorized access, evasion, spam, denial-of-service attacks, or other security abuse.

Simulations, outputs, or automated workflows that are unlawful, unsolicited, misleading, coercive, deceptive, or abusive.

6.2 Deceptive AI, Human Impersonation, or Synthetic Research Misrepresentation

Deploying an AI persona, subagent, simulation, or automated evaluator that falsely claims to be a specific human, licensed professional, government official, emergency responder, employee, customer, expert, public figure, or representative without authorization.

Designing an AI persona, subagent, simulation, or automated evaluator to deny, hide, or misrepresent that it is AI when disclosure is legally required or when the context would make non-disclosure materially deceptive.

Using cloned, synthetic, or similar-sounding voices without valid authorization from the relevant person or rights holder.

Misrepresenting identity, business affiliation, legal authority, research methodology, or the purpose of the interaction.

Creating, presenting, or distributing simulated feedback, AI persona responses, or outputs 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 7.

6.3 Telemarketing, Messaging, and Communications Abuse

At launch, Customers must not request or use tests that contact real people or perform real transactions, including real purchases or payments, and Lythe must not execute those actions. Lythe limits authenticated testing to client-provided test environments and test accounts with safeguards appropriate to the target. Approved test email, if available, is limited to controlled test workflows that do not deliver communications to real-person recipients. Customer must maintain safeguards for test accounts and environments it controls; Lythe must enforce the accepted testing scope and stop affected execution if safeguards fail. Ordinary Lythe account, billing, support, security notices, and customer-requested delivery of run summaries remain distinct from simulated test actions. The Services must not be used for unsolicited outreach, telemarketing, robocalling, or messaging campaigns.

Attempting to bypass or ignore opt-out, suppression, user privacy choices, or data subject rights requirements.

Using the Services to conduct spam campaigns, mass unsolicited outreach, marketing campaigns, or abusive sales or debt collection tactics.

Circumventing platform safety controls, rate limits, monitoring, compliance systems, or other abuse-prevention measures.

6.4 High-Impact and Regulated Decisions

Using Lythe services as the sole basis for any High-Impact Decision, or to make or materially influence a High-Impact Decision without appropriate human review, documentation, compliance controls, and any prior written approval required by Lythe under Section 14.

Representing an AI persona, subagent, simulation, or automated evaluator as a licensed professional, or providing regulated medical, legal, financial, insurance, employment, housing, or public-benefits advice without required authorization, disclosures, compliance controls, and qualified human oversight.

Making eligibility, diagnosis, treatment, credit, insurance, hiring, firing, tenant screening, immigration, or law-enforcement decisions without required compliance controls.

Using Lythe services for emergency dispatch, crisis response, or situations where failure could create immediate risk to life or safety, unless expressly approved in writing by Lythe under a specific regulated deployment agreement.

6.5 Sensitive, Vulnerable, or Protected Groups

Exploiting, manipulating, or unfairly targeting children, older adults, people in crisis, or other individuals in circumstances of vulnerability.

Inferring or using sensitive characteristics such as race, religion, health, disability, sexuality, political opinion, union membership, or biometric identity in ways that violate law or create unfair treatment.

Deploying AI personas, subagents, simulations, or automated evaluators to manipulate political views, suppress voting, mislead voters, impersonate campaigns or officials, or conduct unlawful political outreach.

6.6 Model, System, Platform, and Research Abuse

Attempting to reverse engineer, bypass, overload, degrade, probe, scrape, or misuse Lythe systems or infrastructure.

Using prompts, tool calls, data, or integrations to bypass safety filters, rate limits, access controls, monitoring, compliance controls, or billing systems.

Using Simulation Evaluation or other automated improvement features to coordinate spam, evasion, manipulation, artificial engagement, security abuse, or platform attacks.

Creating benchmark, evaluation, or comparison claims that are misleading, non-reproducible, or falsely attributed to Lythe services.

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

6.7 Unauthorized Target Testing

Users must not direct or allow an AI Persona, subagent, Simulation, automated evaluator, workflow, integration, or Output to access, test, crawl, scrape, scan, monitor, evaluate, interact with, or analyze any target that is not an Authorized Target under the Terms of Service. This prohibition expressly includes competitor and other third-party targets unless the Target Owner gave prior express written authorization covering the testing. Users must not configure agents to conceal testing, evade blocking, bypass controls, exceed approved rates, or continue after any stop signal, withdrawal, expiry, dispute, or scope limit. Before the first run against any target, Customer must submit the Target Certification required by the Terms of Service, identifying the exact target and permitted scope; a material change to the target, environment, method, intensity, or scope requires a new Target Certification. 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.

7. AI Disclosure Requirements

Customers must comply with all applicable AI disclosure and consumer-protection requirements. Where required by law, regulation, platform policy, customer agreement, Lythe instruction, or where non-disclosure would be materially deceptive, customers must clearly disclose that an End User is interacting with an AI system or that feedback, evaluations, personas, or outputs are AI-generated or simulated.

AI-generated simulation outputs should be clearly labeled. For example:

“AI-generated simulation. These findings reflect simulated behavior and stated assumptions; they are not observations of real research participants. Validate consequential findings independently.”

For AI personas, simulations, and automated evaluations, customers must use clear and accurate language when describing outputs, assumptions, limitations, and research methodology. 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.

These examples are provided for illustration only. Customers remain responsible for determining and implementing any additional disclosures, consents, recording notices, opt-out mechanisms, research notices, or human-escalation statements required for their use case and jurisdiction.

8. Call Data Consent and Communications Compliance

Customers are responsible for required notices, consents, lawful bases and other requirements for source recordings, transcripts, conversations, feedback and integration information. Only Permitted Inputs may be supplied or exposed to Lythe. Real customer records, feedback or conversation materials require a separate Data Approval before receipt or access, including approved categories, source, Customer-specific purpose, transfer, access, providers and retention.

Customers must obtain, document, maintain and honor legally required consents and other lawful permissions before collecting or processing source materials. Public availability does not itself establish lawful authority for every use. Minimize identifying material before transfer and complete any required Data Approval; recordings additionally require an enabled feature and specific approval. Real customer records require specific Data Approval before receipt or access,

including the approved source or target and safeguards. No approval permits sensitive or otherwise prohibited inputs.

Customers must comply with applicable privacy and data-protection laws, user notice requirements, opt-out obligations, marketing rules, and regulatory requirements when using the Services.

Customers must promptly honor opt-out or consent revocation requests and maintain suppression records where required.

Customers must provide accurate disclosures, workflow descriptions, research methodology, business identity information, and compliance documentation to Lythe where requested.

Lythe may require proof of consents, data-source documentation, workflow descriptions, research methodology, business identity, opt-out procedures, and compliance controls before enabling or continuing certain use cases.

9. Human Oversight and Escalation

Lythe must provide appropriate human oversight and operational escalation for the AI personas, subagents, simulations, automated evaluators and workflows it operates, including uncertainty, failures, safety concerns and possible activity outside the accepted Run Record. Customers must identify relevant requirements when requesting work, provide an authorized contact for questions and stopping, and apply appropriate human review to their downstream use of reports and Outputs.

AI personas, subagents, simulations, and automated evaluators must have defined authority limits and must not promise outcomes they cannot guarantee.

Lythe must escalate low-confidence results, possible scope breaches and material safety or reliability concerns to an appropriate human reviewer and stop affected execution where required by the accepted safeguards or stop procedure. Customer must promptly report known concerns and may withdraw instructions through the designated stop route. Human review does not permit tests to contact real people or perform real transactions at launch.

Customers must review delivered reports, simulation results, disclosed limitations and escalations as appropriate before relying on them, and promptly report material concerns to Lythe. Lythe remains responsible for reviewing operational logs and supervising its own agents; Customer is responsible for complaints, rights requests and downstream actions within its control, subject to applicable law.

Customers must maintain accountability and appropriate approval controls for permitted integration changes and downstream use of findings. At launch, tests must not contact real people or perform real transactions; human review or a customer request does not create an exception.

10. Customer Content, Prompts, and Knowledge Bases

Customers are responsible for all permitted prompts, scripts, knowledge bases, websites, files, FAQs, permitted CRM-derived datasets, approved conversation-derived information, research data, product flows, screenshots, prototypes, landing pages, user research materials, APIs, and other content they upload, connect, or make available to Lythe services.

Customer content must be accurate, lawful, authorized, and suitable for the intended workflow.

Customers must not supply personal information outside Permitted Inputs. Necessary ordinary personal details incidental to authorized public pages and real customer records or feedback within a Data Approval completed before receipt or access may be processed under the DPA. Sensitive or regulated data, children's data, real payment-card credentials, identifiable raw replay and re-identification keys remain prohibited. 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 and third-party materials require appropriate rights and safeguards; approved test secrets use the separate secure mechanism in the Terms of Service.

Customers must not use third-party content, voices, trademarks, recordings, datasets, websites, product materials, customer records, or personal data without appropriate rights.

Lythe must test the behavior of the AI personas, subagents, simulations and automated evaluators it operates before use and when material changes to prompts, tools, integrations, models, datasets or knowledge sources affect that behavior. Customers must assess delivered Outputs before relying on them and test changes they make to their own systems or downstream workflows.

11. Lythe Evaluation, Simulation, and Automated Improvement Controls

Lythe services may offer Simulation Evaluation, simulation, testing, optimization, monitoring, red-teaming, or automated improvement features. These features are intended to improve the reliability, safety, performance, usability, and compliance of AI personas, subagents, simulations, automated evaluators, product experiences, and related workflows.

Customers may not use Simulation Evaluation or other automated improvement features to evade compliance requirements; manipulate or deceive people; coordinate spam, harassment, synthetic research abuse, abusive communications, unauthorized testing activity, or security abuse; create harmful autonomous behavior; misrepresent synthetic feedback as real research; or deploy materially higher-risk behavior without appropriate controls and approval.

Customers must identify their requested objectives, success criteria, prohibited outcomes and relevant downstream constraints. Before executing accepted work, Lythe must define appropriate

failure conditions, authority limits, escalation thresholds, human-review requirements and stopping or rollback procedures for systems it controls. Customer requests and any optimization remain subject to accepted Run Records, required Run Confirmations and launch restrictions; a request does not authorize automatic deployment of changes to Customer's production systems. Customers remain responsible for approval and rollback controls for downstream changes they make.

Automated improvements must not be deployed in connection with High-Impact Decisions or Regulated Use Cases without appropriate testing, documented review, human approval, ongoing monitoring, and rollback controls.

Customers must ensure source collection, any required anonymization or minimization, Permitted Inputs and downstream use are lawful, authorized, secured and consistent with disclosures and agreements. 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. Simulations must not contact real people or perform real transactions at launch.

Lythe must maintain records reasonably necessary to account for the testing execution it controls, including applicable Run Records, Customer confirmations, evaluation logs, relevant model, prompt and workflow versions, configuration changes, failures and remediation. Customers must retain their authority evidence, submitted requirements, instructions, approvals and records reasonably necessary for their own downstream decisions. Access to retained records remains subject to the applicable agreement and confidentiality and data-protection requirements.

Lythe may investigate, restrict, pause, disable, or require review of Simulation Evaluation, automated optimization, or automated improvement features where Lythe reasonably determines that their use violates this policy or creates legal, safety, security, compliance, reliability, reputational, or platform-integrity risk.

12. Data Privacy and Security

Customers must use Lythe services in compliance with applicable privacy, data protection and security laws. Customers are responsible for required notices, consents and rights requests concerning source materials and downstream use, and for access controls and available retention settings for systems they control. Lythe is responsible for operating its own systems and applying its contractual privacy, security and retention obligations.

Do not submit data that you are not authorized to process.

Do not use the Services to identify, re-identify, profile or target real individuals, or extract personal datasets from public pages. Necessary incidental public-page processing and separately approved feedback may be used only for the permitted evaluation or Customer-specific simulation purpose. Do not reverse anonymization or combine inputs and Outputs to identify people.

Do not connect systems or APIs in a way that exposes data beyond what is necessary for the workflow.

Use appropriate authentication, access controls, environment separation, and internal approval processes.

Promptly notify Lythe of any suspected unauthorized access, misuse, security incident, privacy violation, or harmful agent behavior involving Lythe services or Customer Data.

13. Testing, Monitoring, and Auditability

Lythe must test and supervise the AI personas, subagents, simulations and automated evaluators it operates, with checks appropriate to the accepted scope and risk, including tool use, authority and scope limits, escalation, refusal behavior, edge cases, adversarial inputs and failure recovery. Customers must assess report limitations and research methodology before downstream reliance and maintain appropriate safeguards for their own systems and integrations. Testing remains subject to the Terms of Service's launch restrictions.

Lythe must maintain applicable operational records of its AI persona versions, subagent configurations, prompts, scripts, tools, knowledge sources, simulation settings, evaluation criteria and execution dates in accordance with its contractual retention and data-protection obligations. Customers must retain their submitted requirements, authority evidence, confirmations and records of downstream changes they control.

Lythe must monitor its simulation execution as appropriate to the accepted scope and risk. Customers must review delivered reports and available permitted evidence as appropriate for their intended downstream use and report material concerns; they are not required to supervise Lythe's agents during execution.

Lythe must investigate material abnormal simulation outcomes, unreliable Outputs and actions by its agents outside approved scope. Customers must promptly report known concerns to Lythe and investigate complaints, errors or risks arising from their own systems, materials or downstream use. Each party must cooperate as reasonably required by the applicable agreement.

Where material legal, safety, security, compliance or reliability risk is identified, Customer must promptly notify Lythe and withdraw affected instructions through the designated stop route. Lythe must stop affected agents and simulations and prevent further execution as required by the

accepted stop procedure. Customers must pause or roll back affected integrations, workflows or downstream uses they control.

14. Regulated Use Case Approval

Customers must not deploy Lythe services for a Regulated Use Case without Lythe's prior written approval, unless Lythe expressly states in writing that prior approval is not required for a particular category of use.

Regulated Use Cases include, without limitation, healthcare, financial services, insurance, legal services, employment, housing, education, government services, political outreach, debt collection, emergency response, and workflows involving minors or individuals in circumstances of vulnerability.

Approval may require additional terms, security review, human oversight, audit logs, compliance documentation, business verification, customer prompts, scripts, workflow descriptions, opt-out procedures, research methodology, data retention controls, restrictions on automated decision-making, or other safeguards appropriate to the proposed use.

15. Enforcement

Lythe may investigate potential violations of this policy. Lythe may suspend, limit, throttle, block, disable, or terminate access to the Services, AI personas, subagents, simulations, automated evaluators, APIs, integrations, workflows, uploads, reports, or accounts where Lythe reasonably believes there is a violation, risk of harm, legal exposure, security threat, or platform abuse.

Lythe may also remove content, require remediation, request proof of consent, request proof of rights or data source, notify affected parties where required, cooperate with law enforcement or regulators where legally required, or preserve records for compliance and enforcement purposes.

16. Reporting Abuse

To report suspected misuse of Lythe services, including unlawful uses, impersonation, spam, fraud, unsafe AI behavior, misleading synthetic research, privacy violations, unauthorized third-party website testing or analysis, or security concerns, contact: team@lythe.ai.

17. Relationship to Other Terms

This AI Use Policy supplements the applicable Lythe Terms of Service, Terms of Use, customer agreements, Order Forms, and any applicable product-specific terms that expressly incorporate or reference this policy, including terms for Fraser Simulation Studio.

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.

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 AI Use 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.

18. Changes to This Policy

Lythe may update this AI Use Policy from time to time to reflect product changes, legal developments, AI requirements, model behavior, safety learnings, or operational needs. 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.

19. Contact

Questions about this AI Use Policy may be sent to:

Legal: team@lythe.ai

Abuse reports: team@lythe.ai

Privacy: team@lythe.ai

Security: team@lythe.ai

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