Legal

Locke Privacy Policy

This policy covers **the Locke product** — the desktop app and the browser extension. Our websites are covered separately by the Website Privacy Policy.

Effective: October 17, 2026 · Last Updated: October 17, 2026

The short version. Locke processes requests locally to detect and mask sensitive information before supported AI requests are sent. Outgoing browser requests that carry data, on any website, can be transferred to the Locke desktop app on the same computer so it can tell whether they are AI traffic, and desktop discovery can inspect initial requests to unknown HTTPS destinations, as explained in Section 6. Screening content and detection reports are not uploaded to Sonomos for analysis. Requests that you send go to their intended service, not through our servers. Account services, settings sync, updates, and support or crash reports you choose to send are described in Section 3. Local metadata and crash records have their own retention rules in Section 7.

Which policy do you need? This one covers the **Locke product**. Our Website Privacy Policy covers our websites — sonomos.ai and every sonomos.ai subdomain, including support.sonomos.ai and trust.sonomos.ai. A single set of Terms of Service covers both.

This Locke Privacy Policy for Sonomos, Inc., doing business as Sonomos, a Delaware corporation with its principal place of business at 9924 Kika Court #2416, San Diego, CA 92129 ("Sonomos," "we," "us," or "our"), describes how and why we access, collect, store, use, and/or share ("process") personal information in connection with Sonomos Locke.

If you do not agree with our policies and practices, please do not install or use Locke. If you have questions, contact us at info@sonomos.ai.

1. What This Policy Covers

In Short: The Locke software itself, in every form we ship it.

1.1 The Product. "Locke" means the Sonomos Locke product in each of the forms we distribute it:

- **The Locke desktop application** — including its local daemon, graphical application, command-line tools, installers, updates, and documentation, for Windows and Linux, with macOS planned.
- **The Locke browser extension** — including its detection and masking features, together with related local software components, on every browser we support.

1.2 Why This Policy Is Separate. Our websites and Locke are different kinds of thing, and treating them under one privacy policy would make that policy less accurate about both. A website processes your data on a server. Locke's screening pipeline processes content on your computer; account services and information you choose to send to support are separate. Those facts deserve to be stated plainly rather than averaged together, so we publish them separately.

1.3 What This Policy Does Not Cover. This policy does not cover our websites — sonomos.ai and its subdomains — which are governed by the Website Privacy Policy at sonomos.ai/privacy. It also does not cover the third-party AI tools and services you use Locke with, such as ChatGPT, Claude, Gemini, or a locally hosted model. Once you choose to send a message, the receiving service handles it under its own privacy policy. Locke's role is to determine what is in that message before you send it.

2. What Stays on Your Device

In Short: Screening runs on your device, including the transfer from the extension to the desktop app.

The screening activities below happen on your device. Request content is not sent to Sonomos for screening. The request may then be sent to its intended destination, allowed unchanged or with redactions, or blocked. Support and crash information that you explicitly send to us is a separate flow described in Section 3.

- **The content you type, paste, dictate, or scan.** Locke inspects it in memory on your machine.
- **Files and attachments you scan.** Document parsing and inspection is performed locally.
- **Detection.** Sensitive-data detection uses on-device pattern matching and on-device machine-learning models. No cloud model, inference API, or remote classifier is involved.
- **Masking.** Sensitive values are replaced locally, before the message leaves your device.
- **What was detected.** The categories, counts, positions, and values of anything Locke finds stay on your device. Detection results are not uploaded through the screening or account/settings pipeline.
- **Detection metadata.** Local records can include the detection type, timestamp, match length, direction, engine and confidence, provider, redaction status, source, capture surface and media type, together with a session-scoped value digest used to recognise repeated values. Session records include app and engine versions, operating-system information, subscription tier and session times. Some processing paths also support structural field paths. These records do not store the matched text itself, but metadata and digests can still reveal information about your activity and should be treated as private.
- **Desktop discovery records.** The local proxy keeps per-install hashed host history and classification state. Separate local discovery diagnostics can contain hostnames and process identifiers, classification scores and structural request metadata. Structural field names and paths can themselves be identifying even when content values are excluded.
- **Extension state.** The browser stores disabled-site choices and a rolling 100-entry diagnostic audit covering connection and screening status, policy-key names, errors and CSP violations. These records can include URLs and source paths; captured request bodies are not stored. Connection status and screening counters are also kept for the browser session.
- **Your reports.** Compliance and activity reports are generated locally from local data and are yours to export or delete. We do not receive them.

We do not receive the screening content or local detection database through the screening pipeline. That is different from saying Locke processes no data: the local handling described here is what makes the protection work.

3. What Locke Sends to Our Servers

In Short: Account, settings and update services, plus support information only when you send it.

Locke uses the following remote services, separately from local screening:

- **License validation and subscription status.** Locke checks that your subscription is active. The desktop application repeats this check periodically while it is running. The request carries your account identifier and subscription tier, and no content, detection data, or usage data. Like any network request, it discloses your IP address and the fact that Locke is running to us.
- **Account authentication.** Locke requires a Sonomos account. When you sign in, Locke authenticates through our authentication provider, Clerk, which holds your name, email address, and session activity. Signing in is required to use Locke.
- **Your settings.** The settings you configure — including which AI services you have asked Locke to monitor and which categories of sensitive data you have enabled — are synced to your account so that they follow you across devices and, on team plans, so that your administrator can apply a policy. This is a record of your configuration choices. It is not a record of the sites you visit, and it carries none of your content.
- **Team and account actions.** Where available, the app exchanges team membership, invitations, policy and billing information, and requests approval for controlled deletion operations. The deletion-approval flow does not send detection content, categories or counts.
- **Update checks and delivery.** The desktop application checks for updates through our designated release channel. The browser extension is updated automatically through the applicable browser marketplace, such as the Chrome Web Store, under that marketplace's own terms.

Support messages and crash reports. The desktop app automatically saves up to three local crash records. They can contain the time, error message and stack trace, app and Electron versions, platform and additional diagnostic details. The app attempts to remove home/profile paths and email-like strings and limits the size of these records; this does not guarantee that all sensitive information is removed from an error message. Saving a crash locally is separate from sending it: a crash report or support message is transmitted to Sonomos only when you choose to send it. Review it first and do not add confidential content. A submission can include your written message or note and a reply email address if you provide one.

Certificate checks. Certificate validation can also make standard certificate-revocation requests. The extension separately makes local connection and presence checks on your own computer, described in Section 6.

Requests you direct. When you send a prompt to an AI provider, that request goes to the provider you chose, at your direction. We are not the recipient of it and it does not pass through our servers.

4. What Locke Never Does

Locke does not:

- Upload your request bodies, files or local detection database to Sonomos for screening.
- Sync raw detected values, detection categories or counts through its account/settings services.
- Use cloud-based AI or machine-learning services for content analysis.
- Build a general browsing-history or keystroke log for Sonomos. Local request processing, discovery and operational metadata are described in Sections 2 and 6.
- Use cookies, web beacons, pixels, fingerprinting, or any tracking mechanism.
- Send background behavioural analytics to Sonomos. Local detection and operational records exist, as do locally captured crash reports; crash and support submissions only travel when you send them, as described in Section 3.
- Sell or rent your data, or use it for advertising or profiling.
- Use your data to train our models or anyone else's.

Local screening reports are not synced as usage analytics to Sonomos. What we can see is that a licence check happened for your account, and what settings you chose. We do not build a usage history out of that, and we retain only your current subscription status — but we would rather tell you that is a commitment we are making than let you believe the architecture makes it impossible.

5. Personal Information We Hold Because of Locke

In Short: Account and billing records, synced configuration, and information you choose to send us.

Using Locke requires a Sonomos account, and a paid tier requires a subscription. The personal information behind those — your name, email address, account identifier, and subscription and billing records — is described in full in the Website Privacy Policy, including how long it is kept and how to have it deleted. Locke's own account and settings data is held by our authentication provider, Clerk, and in our hosted database, provided by Supabase; billing runs through Stripe.

Beyond the account record, we hold the synced settings, team/account-action information and any support or crash submissions described in Section 3.

In the categories used by California law, in the preceding 12 months Locke has caused us to process: **A. Identifiers** (account identifier and email address, for authentication and licence validation), and **F. Internet or other electronic network activity**, limited to the IP address and connection metadata that any network request necessarily discloses, and to the configuration you have chosen. These categories describe the routine remote account flow, not every type of information Locke may process locally or that you may choose to include in a support message. Do not include sensitive personal information in support submissions.

6. Permissions Locke Requests

In Short: Browser request access and desktop network access are different capabilities.

6.1 Browser extension. The current extension loads on every http and https page you visit, as declared in its manifest. It does not load on local files. This is separate from desktop discovery below. Its page hooks process outbound request bodies, including supported files and attachments, and request metadata such as the method, destination host, path/query and page-set headers. Browser-added cookies are not part of this capture, and the delegated-upload path strips presigned query credentials.

Classification of every request. Locke does not decide in the browser which requests are AI traffic. For fetch/XHR requests with a body, on any website, the extension holds the request and transfers it through browser native messaging to the Locke desktop app on the same computer. There, a local classifier first decides whether the request looks like AI traffic. Requests it classifies as not AI are released unchanged without being screened for sensitive information. Requests it classifies as AI, and requests to the paths of known AI services that are listed in the extension's built-in catalog, are screened, and the app's allow, redact or block verdict is applied. The extension tells the app whether a request is to a known AI service's path. Request content is not sent to Sonomos. This local transfer is data processing even though Sonomos does not receive the content.

Certain attachment writes to third-party object storage are always screened when initiated by a page of a known AI service. The normal request still goes to its intended destination if allowed. If the desktop app cannot be reached, requests are blocked rather than sent unchecked, unless you have turned on a time-limited fail-open window in the app's settings. One exception applies when you have deliberately signed out of the Locke desktop app (and your device is not under an organisation lock): requests to ordinary websites are then sent unchanged and unclassified, so that sign-in pages and forms keep working, while requests to known AI services, from pages of known AI services, to local or private-network addresses, and file uploads are still blocked. Beacons and other requests that cannot wait for a verdict are blocked on the paths of known AI services and otherwise left alone, because they cannot be classified. Where the browser asks for your consent before the extension may share page data with the desktop app, nothing is transferred until you give it; until then, requests to known AI services are blocked and other requests are left alone.

Classifying every request is not a guarantee that every AI request is screened: the classifier can be wrong, and navigation/query-string prompts, WebSockets, worker-originated traffic and unsupported upload forms have coverage gaps. Per-site settings can disable screening. The extension page and coverage notes explain the limits. Native-host checks and a loopback presence signal also exchange browser/version information locally; Chromium registration additionally sends the extension ID.

6.2 Desktop application. When desktop capture is active, supported HTTPS traffic is routed through a local proxy using the platform's proxy/routing and certificate-trust configuration. The current Locke configuration enables active AI-service discovery. For eligible previously unrecognised HTTPS destinations that reach that

proxy, discovery may decrypt and inspect initial requests before knowing whether they are AI, including services later determined not to be AI. Its local classification receives the destination, method, path/query, headers and decoded body, then serializes structural features while stripping content values from that representation. Field names and paths can still be identifying. Requests identified as AI are handled under the applicable screening and organisation settings.

Discovery has request-count and time limits for its classification window, but these are not an absolute limit on traffic decrypted through an already-open intercepted connection. Applicable host allow/block rules, excluded infrastructure and unsupported connection types affect which connections are inspected. Coverage depends on the platform and application: for example, the Linux capture shim bypasses positively identified extension-covered browsers, and applications that bypass the configured proxy or reject interception may not be covered. AI-service action controls apply after classification; changing those controls does not by itself mean initial discovery inspection is disabled.

6.3 Local storage. Locke keeps local settings, detection metadata and reports, extension state, operational logs and crash records. These have different deletion paths; clearing detection history does not mean every local record is erased. Uninstalling does not necessarily delete retained user data, and the browser extension is removed separately. See Section 7 and the uninstall guide.

7. Retention

Screening content — Request bodies and files are processed transiently for local screening and are not stored as raw content in the detection database or uploaded to us by that pipeline. This does not cover files you already saved, your exports, or information that may appear in local error records or that you explicitly include in a support submission.

Detection results and reports — Kept locally indefinitely by default, with automatic age-based deletion off. You can request deletion through Locke, subject to any applicable team approval, or configure supported retention controls. A changed automatic-retention policy takes effect when the metrics service restarts. Exported copies remain wherever you saved them.

Deletion audit — A separate local deletion audit survives the detection database's clear-all operation. It records the operation time, requesting interface, scope, affected counts and errors. Quarantine records can include filesystem paths. Clearing detection history is not deletion of this audit.

Desktop discovery records — Host history is limited to 4,096 entries. Confirmed AI entries have a 30-day lifetime refreshed on use; suspected, non-AI and unscreenable states generally expire after 7 days, while pinned/mTLS states use 30 days. Separate discovery diagnostics rotate at 32 MiB per file with at most two generations; this is a size limit, not an age-based deletion schedule. These records are separate from the detection database and its clear-all operation.

Extension state and crash records — The extension's persistent operational audit rolls over after 100 entries; disabled-site choices remain until changed or the extension's stored data is removed. Browser-session status and counters have session storage lifetime. The desktop app keeps up to three local crash records, removing a report after a successful submission or when you discard it. Remove retained local data separately when uninstalling if you want it erased.

Your synced settings — Held for as long as your account is open, and deleted with it.

Account and billing records — Retained as set out in the Website Privacy Policy, which explains why billing records are kept longer than the rest.

License-validation requests — Not retained as a usage history. We hold your current subscription status, not a log of when you used Locke.

8. Security

- **Local-first architecture** — The strongest control here is structural: content that never leaves your device cannot be intercepted, breached, or compelled from us.
- **Encryption in transit** — TLS for license validation, authentication, and update delivery.
- **Signed releases** — Desktop releases are signed, and browser-extension updates are distributed through the applicable browser marketplace.
- **Least privilege** — Browser permissions and desktop interception capabilities are described separately in Section 6; supported coverage and bypasses depend on the installed version and platform.

What local processing does not protect against. Running on your device means Locke's protection is bounded by your device's own security. It cannot defend against malware, keyloggers, screen capture, unauthorized physical access, or a compromised operating system. It also cannot recover data you have already sent to a third party. More detail on our threat model is on the Locke security page.

Detection is not perfect. Locke uses pattern matching and on-device machine learning, and can produce both false positives and false negatives. It is a tool to support your judgment, not a substitute for it, and using it does not by itself constitute compliance with HIPAA, GDPR, GLBA, or any other framework.

9. Children

Locke is not directed to, and may not be used by, anyone under 18. We do not knowingly collect personal information from children under 18. If we learn that we have, we will deactivate the account and delete the data. Contact info@sonomos.ai if you become aware of any such collection.

10. Your Privacy Rights

The rights described in the Website Privacy Policy — including the CCPA/CPRA rights of California residents, the rights available under other US state privacy laws, and GDPR, UK GDPR, PIPEDA, and Australian and New Zealand equivalents — apply in full to the personal information we hold in connection with Locke.

You do not need to submit separate requests for our websites and for the account and support information we hold because of Locke. One request covers both. To exercise a right, visit sonomos.ai/contact, email info@sonomos.ai, or write to Sonomos, Inc., 9924 Kika Court #2416, San Diego, CA 92129. We respond to verified requests within 45 days.

We cannot retrieve or erase content processed only locally and not submitted to us. Local data is under your control, subject to applicable team controls and the deletion exceptions in Section 7; uninstalling alone may retain it.

11. International Users

Our account, licensing, and billing infrastructure is hosted in the United States. If you use Locke from outside the United States, the account information described in Section 5 will be transferred to and processed in the United States, under Standard Contractual Clauses and other appropriate safeguards where required.

Screening content is processed on your device, including browser-to-desktop transfer, rather than sent to Sonomos for analysis. Requests you send to third-party services and support information you choose to submit can be transferred to their respective recipients.

12. Verifying These Claims

A privacy product should not ask for faith. Our Trust & Transparency page sets out what actually leaves your device in checkable terms, along with an honest account of what is not yet true, and our Locke security page describes the architecture in more detail. If you find a discrepancy between this policy and Locke's behavior, tell us at info@sonomos.ai — we would rather fix it than defend it.

13. Changes to This Policy

We may update this policy from time to time. Material changes will be indicated by an updated "Last Updated" date above and may include direct notification. If a future release of Locke changes what it sends to our servers, we will update this policy before that release ships, not after.

14. Contact Us

Sonomos, Inc.
9924 Kika Court #2416
San Diego, CA 92129
United States

General: info@sonomos.ai
Online: sonomos.ai/contact