

Privacy Policy - Nyx AI

Last updated: 20 August 2026 **Version:** 1.42 **Applies to:** Nyx AI desktop application v1.4.x and later (distributed under the name "Nyxion" before version 1.5.438; both names refer to the same software, and existing on-disk folder names such as `Documents\Nyxion` and the app's user-data folder keep the old name so no data moves)

1. Summary

Nyx AI is a local-first desktop application. Nyx AI itself does not operate user accounts, a hosted backend, telemetry, analytics, crash reporting, payment processing, or advertising, and it sends us no usage or diagnostic data of its own.

In the current product, the personal data we receive directly is the correspondence you choose to send to customerservice@nyxlimited.com, plus the IP address and app version carried by the automatic update check (on by default, opt-out — see [Section 6.4](#)). Separately, our public **website (nyxai.uk)** is live; visiting it sends the IP address and standard HTTPS request/security metadata needed for Cloudflare to host and protect it. The Website includes optional Pro-waitlist and one-time-support-payment flows; those flows operate only where they are configured and offered. If you choose to use them, the processing by Brevo, Stripe, Cloudflare Turnstile, and the configured rate-limit service described in [Section 6.5](#) applies. The website does not currently use analytics and does not currently offer a public app installer. The desktop application performs none of that website processing. Your code, files, prompts, chat history, API keys, settings, local indexes, audit logs, and generated outputs are stored on your machine; they are transmitted only when a network feature needs them — when you choose to use a third-party service, and through the automatic requests that enabled or default-enabled features make (such as the update check and the flows in [Section 7](#)). On-device features such as Read aloud (text-to-speech) run entirely on your machine by calling Chromium's Web Speech API and selecting only your operating system's installed offline voices: no text or audio is transmitted. Nyx AI does not download, install, bundle, repackage, or claim ownership of Microsoft or other platform voices or a speech engine; those remain provided by their publishers under their own terms.

One distinction matters for how you obtained Nyx AI. Aside from an automatic update check — on by default, which you can turn off, and which carries your app version plus the IP address and standard request metadata any internet request reveals (see [Section 6.4](#)) — the application does not transmit anything about you or your use to NYX LIMITED. Separately, the platform that delivers and runs it can: if you install Nyx AI from the **Microsoft Store**, Microsoft (as the store operator) gives us mostly aggregate install/usage analytics and may provide event-level Health diagnostics such as an individual stack trace or crash dump, and Windows (as the operating system) collects diagnostic data, each under Microsoft's own terms rather than ours. We describe this in [Section 6](#). In short: Nyx AI sends no project content, prompts, usage telemetry, or diagnostic data to us on its own; the

automatic/manual update check sends the app version — plus the IP address and standard request metadata any internet request reveals — to our documented update host.

When you use a third-party AI model provider, MCP server, Google Drive integration, website fetch/search tool, marketplace source or plugin repository, language-server download, cloud vision handoff, an automatic background helper model call (file summaries, multi-file refactor checks, memory reflection), or similar integration, data may go directly from your machine to that third party. Those providers are not operated by us and their own terms and privacy notices apply.

2. Who We Are

Nyx AI is provided by **NYX LIMITED**, a company registered in England and Wales under company number 17261427, whose registered office is at 70 Clarkehouse Road, Sheffield, England, S10 2LJ ("we", "us", "our"). Contact: customerservice@nyxlimited.com.

For the personal data we receive in support, legal, security, or other correspondence, and for the website waitlist, payment-administration, and abuse-prevention processing described below where those features are offered, NYX LIMITED is the controller. The same applies to the Microsoft Store developer reports Microsoft supplies to us for Store installs ([Section 6.1](#)): to the extent event-level Health data (such as a stack trace or crash-dump file) contains personal data, NYX LIMITED is the controller of the copy we receive, while Microsoft's own collection remains governed by Microsoft's terms. We are not the controller or processor for your local project data that never reaches us.

If you configure Nyx AI to send prompts, code, files, embeddings, messages, or other data to a third-party provider, your relationship with that provider is separate. Depending on the provider and your use case, that provider may act as an independent controller, a processor for you, or another legal role under its own terms.

3. Scope

This policy covers:

- The Nyx AI desktop application
- Our public website, **nyxai.uk**, from the time it is reachable ([Section 6.5](#))
- Support, legal, security, accessibility, and other correspondence you send us
- Current local-first functionality and optional integrations described in [THREAT-MODEL.md](#)

This policy does not cover:

- Third-party AI providers
- Third-party MCP servers, plugins, hooks, package registries, marketplaces, or websites
- Google Drive, Telegram, GitHub, Ollama, OpenAI, Anthropic, Google, or other providers you choose to use
- Your own project code, services, websites, or users

4. What We Collect

4.1 Correspondence you send us

If you email us, we receive the information contained in that email, which may include:

- Your name
- Your email address
- Message content
- Attachments or logs you choose to include
- Technical details you choose to provide, such as version, platform, or reproduction steps

We use this information to respond to you, investigate the issue, keep appropriate records, and protect our legal and security interests.

Please do not include health or disability diagnoses, allegations or records of criminal offences, or other highly sensitive information unless it is strictly necessary for us to handle your request. If we receive such information, we minimise or delete it unless retaining it is necessary and permitted under a documented additional condition in UK data-protection law.

AI-generated content reports. Nyx AI includes an in-app way to report an AI-generated response. The report is prepared locally and shown to you for review before Nyx AI does anything else with it. It contains only the reported AI response, your reason, the app version, and the active model name. The response and your reason are passed through Nyx AI's privacy sanitizer before the report is shown, copied, or placed into an email draft; workspace-root paths are collapsed where possible; and long text is capped. It does **not** include your activity log, workspace location, file paths or URLs beyond any sanitized text you choose to report, operating-system details, or an install identifier. Nyx AI never sends the report automatically, and we receive it only if you press Send in your email application. Stated precisely, because "nothing leaves your machine" would be too strong: **Copy** places the report on your system clipboard and Nyx AI sends it nowhere — though if you have enabled your operating system's cloud clipboard or clipboard-history sync, your OS may itself copy clipboard contents off the device under its own terms. Choosing **Open draft** hands the prepared text to whichever application your system has registered for email. If that handler is a webmail site or a cloud-synced client, your email provider may receive, sync, or store the draft at that point — before you press Send, and under that provider's own terms rather than this policy. NYX itself receives nothing until you send. A report you send us is handled as correspondence under this policy.

Bug, security, and general contact drafts. The in-app "Report a bug", "Report a security issue", and contact buttons open ordinary email drafts. They do not attach logs, prompts, files, activity history, model transcripts, or an install identifier. Some drafts may pre-fill a short subject or template, such as the app version in a security-report template, but you decide what to send.

Diagnostics and reports you copy yourself. Some screens let you copy diagnostic text or a report to your clipboard (for example, the "Copy report" button on an AI-content report). If you copy

that text and paste it somewhere else, that is your own action: where it then goes is outside our control and is not a transmission Nyx AI makes to us.

4.2 Data Nyx AI does not send to us

Nyx AI does not send us:

- Source code, repository contents, file names, or file paths
- Prompts or conversations
- API keys, tokens, passwords, or credentials
- Chat history, local memory, embeddings, semantic index data, or audit logs
- Shell command history or output
- Preview browser history
- Telemetry, analytics, crash reports, usage metrics, or install/device identifiers transmitted to us by Nyx AI (a random install id and a usage counter are stored locally on your device and are never sent to us)

An AI-generated content report that you choose to email us yourself may include the sanitized response, your sanitized reason, the app version, and the active model name described in section 4.1. A bug, security, or contact email may include whatever you type or leave in the email draft. Nyx AI itself does not transmit these to us.

Separately, the automatic update check (section 6.4) contacts an update host we operate and reveals your IP address and app version in the ordinary way any web request does. It sends none of the content or identifiers listed above — no code, prompts, credentials, history, or install/device identifiers.

This describes what the Nyx AI application does. It does not describe what the **platform** that delivers and runs Nyx AI may collect. If you installed Nyx AI from the Microsoft Store, Microsoft provides mostly aggregate analytics and may provide event-level Health diagnostics, and Windows collects operating-system diagnostic data, independently of the application. That platform-level processing is governed by Microsoft's own terms and is described in [Section 6](#).

4.3 Microsoft Store developer reports (Store installs only)

If you installed Nyx AI from the Microsoft Store, Microsoft makes developer reports available to us as described in [Section 6.1](#). These are mostly aggregate acquisition, usage, and ratings analytics, but the **Health** reports may sometimes include **event-level data such as an individual stack trace or crash-dump file**. We use this data for reliability, debugging, and product security — to understand and fix crashes, hangs, and memory failures. We do not use it to identify you, and we do not combine it with your correspondence unless you contact us about the same issue. Any event-level Health data we download from Microsoft's reporting tools is retained only while the referenced defect or reliability issue is being investigated, then deleted; the underlying reports remain available to us within Microsoft's developer tools under Microsoft's own retention.

4.4 Optional website waitlist and support payments

Pro launch waitlist. Where the waitlist is offered, we collect the email address you submit. Brevo sends the confirmation email and keeps the double-opt-in record. Your address is added only after you use that confirmation link. We use it solely for a one-off notification when Nyx AI Pro launches. You can withdraw at any time through the unsubscribe link or by contacting us.

One-time support payments. Where the website offers a voluntary support payment, Stripe processes the payment. We may receive your billing name and email address, transaction outcome, and limited card metadata such as the last four digits. Your full card number and CVC are entered with Stripe and never reach our servers. A support payment does not buy goods, services, a licence, preferential treatment, or product access.

Form and abuse prevention. Website forms may use Cloudflare Turnstile, which processes IP, device/browser, and interaction signals to distinguish users from automated abuse. Where configured, the website also uses Upstash Redis for short-lived rate-limit and replay-prevention keys. Raw email addresses and raw Turnstile tokens are not stored in those keys; the website uses short-lived hashed markers as described in Sections 6.5 and 11.

5. Lawful Bases

We rely on the following lawful bases under the UK GDPR and, where applicable, the EU GDPR:

Purpose	Personal data	Lawful basis
Responding to support, legal, security, accessibility, or other correspondence	Name, email address, message content, attachments you provide	Legitimate interests: responding to correspondence, supporting the product, improving safety, and keeping records
Handling a data-protection complaint or rights request	Identity/contact details, request content, correspondence	Legal obligation where the law requires us to respond; legitimate interests for record-keeping and dispute handling
Establishing, exercising, or defending legal claims	Relevant correspondence and records	Legitimate interests; legal obligation where applicable
Checking for and delivering application updates (on by default; you can turn it off in Settings → Advanced → Updates)	IP address (inherent to any HTTPS request) and app version	Legitimate interests: keeping the application current and secure and delivering updates and security fixes
Reviewing Microsoft Store developer reports for Store installs (Section 6.1)	Mostly aggregate acquisition/usage/ratings analytics; event-level Health data Microsoft may supply, such as a stack trace or crash-dump file	Legitimate interests: reliability, debugging, and product security

Purpose	Personal data	Lawful basis
Sending a one-off Nyx AI Pro launch notification, where the waitlist is offered	Email address and double-opt-in record	Consent, given by signing up and confirming; you may withdraw it at any time
Operating our website and protecting forms from spam and abuse	IP address, standard HTTPS request/security metadata, device/browser and interaction signals, and short-lived rate-limit/replay markers	Legitimate interests: keeping the website available, secure, and abuse-free
Administering voluntary website support payments, where offered	Billing name and email, transaction record, and limited card metadata from Stripe	Legitimate interests: receiving and administering voluntary support; legal obligation for tax and accounting records

We do not operate website analytics, advertising, a general marketing list, or cross-site profiling. The Pro waitlist, where offered, is limited to the one-off launch notification described above. We do not sell personal data, share personal data for cross-context behavioural advertising, or process special category data intentionally.

If we introduce accounts, optional telemetry, crash reporting, marketing beyond the described one-off waitlist notice, hosted services, paid app features, or business support in the future, we will update this policy before that processing begins and obtain consent where consent is required.

6. How You Obtained Nyx AI: Store Distribution and Windows Diagnostics

The public website offers no direct app download. Nyx AI is distributed through the **Microsoft Store as an unpackaged Win32 (EXE/MSI) application** and may also be distributed through another source we expressly operate or designate. Nyx AI itself sends no telemetry, analytics, or crash data in either case. However, the channel through which you obtain it, and the operating system you run it on, may involve data processing **by Microsoft, under Microsoft's own terms rather than ours**. We do not control that processing and do not receive any data Microsoft does not choose to give us.

6.1 Microsoft Store installs

If you install Nyx AI from the Microsoft Store, **Microsoft** (as the store operator) collects data about the app and provides us with **mostly aggregate analytics plus some event-level Health diagnostics**. Based on Microsoft's published developer documentation, this typically covers:

- **Acquisitions and installs** (for example, how many times the app was acquired or installed).
- **Usage**, reported as **active devices and average engagement minutes** — that is, device- and time-based measures, not counts of unique, identified users.
- **Health**, such as crashes, hangs, and memory failures, which **may sometimes include stack traces or crash-dump files**.

- **Ratings and reviews** you choose to leave in the Store.

Microsoft determines which data appears in developer reports according to users' Windows and Microsoft diagnostic settings and Microsoft's current reporting rules. Most of what we see is aggregate (for example active-device and engagement totals) and is intended to tell us roughly how the app is performing, not to identify you. The exception is the Health data above: a crash report can be event-level — an individual stack trace or crash-dump file for a single crash — rather than an aggregate count.

The precise data Microsoft collects is governed by Microsoft's own terms, not ours.

Microsoft provides aggregate install/usage reporting and can provide event-level crash diagnostics; the exact fields, granularity, and retention are determined by Microsoft under the [Microsoft Privacy Statement](#) and the [Microsoft Services Agreement](#) (including its Standard Application License Terms for Store apps). This processing sits outside Nyx AI and outside our control, and we use the reports only to understand reach and reliability and to improve the product.

6.2 Windows operating-system diagnostics

Separately from the Store, **Windows itself** collects diagnostic data about your device and the software running on it. This applies to **all installs of Nyx AI, including a direct download**, because it is a function of the operating system rather than of the application or the Store:

- **Required diagnostic data** is collected by Windows at all times to keep the device secure and up to date.
- **Optional diagnostic data** — which can include app-activity information, crash dumps, and Windows Error Reporting data — is collected according to **your Windows "Diagnostic data" setting** (Settings, then Privacy & security, then Diagnostics & feedback).

This diagnostic data is collected and processed by **Microsoft under the [Microsoft Privacy Statement](#)**, and you control the optional portion through your Windows settings. It is not sent to us, and it is not something Nyx AI enables, configures, or can switch off on your behalf.

6.3 Direct download (non-Store)

If you obtain Nyx AI as a direct download rather than from the Microsoft Store, **none of the Store analytics in section 6.1 apply** — there is no Store acquisition, usage, or health reporting for that copy. Only the **Windows operating-system diagnostics described in section 6.2** apply, because those come from the operating system regardless of how the application was installed.

6.4 Update checks and auto-update

Nyx AI includes an automatic update **check**. It is **on by default**, and you can turn it off in the app under **Settings** → **Advanced** → **Updates** (after which updates are fully manual). When it is on, Nyx AI checks shortly after launch and about once a day while running.

What the check does. It makes a single HTTPS request to an update host we operate (served via Cloudflare) to read a small version file. If a newer version exists, Nyx AI tells you. It **never**

downloads or installs an update without you choosing to — downloading and installing are separate, explicit actions you take.

What the check sends. The update server sees your device's **IP address** and the **app version** (carried in the request's User-Agent), as any ordinary web request would. An IP address can itself be personal data, so we do not claim the request is anonymous; what it does **not** carry is any account, usage data, file or prompt content, install ID, or device identifier — not your code, files, prompts, chat history, or settings.

What can be seen and stored. Because the request travels over the internet to a host we operate through **Cloudflare**, the request's IP address, time, and the version requested appear in Cloudflare's ordinary server/CDN access logs. NYX LIMITED does not separately store those logs; Cloudflare retains them under its own documented retention for the services we use. We do not use these requests to build profiles of you or to identify you, and we keep no application-level record of who checked for updates (see [Section 11](#)). On the in-app updater path, the installer and redirects must remain on the same host; before Nyx AI can run it, its code signature and checksum are verified and its signed embedded version must match the update manifest. If Nyx AI instead opens a manual download or release-information page in your browser, that browser download and any later manual installation are outside the in-app verification path; check the published signature and checksum before running it.

Turning it off. Switching the setting off stops the automatic check entirely; Nyx AI will then contact the update host only if you click "Check for updates" yourself.

Operating-system update behaviour is unaffected: because the Microsoft Store does not automatically update unpackaged EXE/MSI apps, keeping Nyx AI current relies on this check or on you installing a newer build manually.

6.5 Our website (nyxai.uk)

We operate a public website at <https://nyxai.uk> (the "Website"). It provides product and legal information and includes an optional Pro launch waitlist and optional one-time support-payment flow; each operates only where its service credentials are configured and the feature is offered. The Website does not currently use analytics and, as of 20 August 2026, does not offer a public app installer. The Website is separate from the desktop application: nothing in this section is collected by or through the Nyx AI app itself, and installing the app does not by itself involve the processing below.

When you visit the Website:

- **Hosting, CDN, and security (Cloudflare).** The Website is served through Cloudflare, which necessarily processes your IP address and standard HTTPS request/security metadata and may apply DDoS and bot filtering as our infrastructure provider. NYX LIMITED does not separately copy Cloudflare's access logs into its own store.
- **No analytics or advertising cookies.** The Website does not currently run analytics of any kind. It does not load Google Analytics or any other analytics or advertising service, and sets no

analytics, advertising, preference, language, theme, or consent-choice cookies. We may in future introduce cookieless, aggregate measurement that sets no cookie and builds no profile. If we do, we will update this policy to say what it collects, who processes it, and how to object, **before** that measurement begins — you will not find analytics running here undisclosed.

- **Optional Pro waitlist (Brevo).** Where offered, the Website sends the email address you submit to Brevo for a double-opt-in confirmation. The address is added only after you confirm, is used only for one launch notification, and can be removed through the unsubscribe link or by contacting us.
- **Optional one-time support payment (Stripe).** Where offered and selected, Stripe receives the payment and card details directly. We receive only the transaction and limited billing/card information described in Section 4.4. No goods, services, licence, or access are supplied in return.
- **Form protection and rate limiting.** Where a form is offered, Cloudflare Turnstile may process IP, device/browser, and interaction signals. When the shared rate-limit store is configured, Upstash Redis holds short-lived IP/prefix keys and SHA-256 markers derived from an email address or Turnstile token. Raw email addresses and raw Turnstile tokens are not stored in those keys; the markers expire automatically, no later than 24 hours after their relevant protection window. If the shared store is unavailable, per-process limits continue but shared replay protection may be reduced.
- **No public installer.** The Website does not currently provide a public Nyx AI installer. Distribution through Microsoft Store or another source expressly operated or designated by NYX LIMITED is described in Section 6.

The Website's lack of analytics and advertising cookies is described in [Section 8](#), the lawful bases in [Section 5](#), providers and transfers in [Sections 9 and 10](#), and retention in [Section 11](#).

7. Optional Third-Party Services

Nyx AI can transmit data to third parties only when you configure, trigger, or enable features that need them. Several disclosed flows can then operate automatically — including the update check, scheduled work you create, provider/catalog checks, semantic embeddings, chat titles, and background model helpers. The examples below are summarized in the more detailed network-egress table in the Threat Model:

- **AI model providers:** prompts, context, selected files or snippets, and image attachments may be sent to the provider you choose, or to the vision provider resolved by the application's vision settings. This covers images you explicitly ask Nyx AI to analyse **and** screenshots Nyx AI captures itself during preview or build work in order to describe them back to the model.
- **Local Ollama:** data stays on the device only when the resolved Ollama endpoint is loopback and the selected model is not cloud-proxied. A remote endpoint or cloud-proxied Ollama model can send model work off the device and is treated as cloud by Nyx AI's egress controls.
- **Cloud embeddings:** code chunks, file/symbol metadata, and search queries may be sent to a configured cloud embedding endpoint when embeddings are set to Cloud, or when Auto mode is

active while you are using a cloud chat model and a cloud embedding endpoint/key is configured. Local mode and Off mode do not send code to cloud embeddings.

- **Background helper tasks:** to support certain features, Nyx AI may make automatic model calls — for example file summaries, multi-file refactor checks, memory reflection, pre-edit safety checks, classifying a prompt, and naming or titling a chat. Small helper roles prefer a reachable local model; heavier planner/coder/reviewer/reasoning roles follow your model routing and may use an eligible configured cloud provider. Prefer-local routing favours a suitable local model when one is installed, reachable and eligible, but may fall through to an eligible cloud provider when none is available; it is a preference, not an egress block. A model you explicitly choose still wins, subject to the documented same-provider availability fallback. The entire cloud-bound model request — the current message, prior chat history, compacted summaries, memory snippets, file context, and tool results — is passed through Nyx AI's privacy sanitizer before transport; image attachments' pixel data is not text and is not processed by it, and MCP-server and web-tool traffic is separate and is not passed through it (text sent to a cloud embeddings endpoint is scrubbed separately). By default, the sanitizer targets recognised credentials and secrets only; ordinary personal information such as email addresses, phone numbers, file paths, and IP addresses remains unless you select a stronger preset. Sanitization is best-effort and cannot guarantee removal of every sensitive value. These calls add no destination beyond the AI provider you configured and send nothing to us. The free approval-card command explainer is presentation-only and sends neither the command nor a model request.
- **Scheduled tasks:** tasks you create may run automatically and may send data to the provider, service, or destination configured in that task.
- **Web fetch/search:** URLs, queries, and fetched pages may be requested when you or the model trigger those tools under Nyx AI's guardrails. Browser-rendered search uses an ephemeral managed browser profile; a search provider may set cookies or similar state inside that short-lived profile, but Nyx AI does not click the provider's consent/accept buttons on your behalf.
- **Document image assistance:** document writers may fetch remote images named in document arguments or Markdown/HTML when the document-image and web-access controls allow it. Presentation generation can also derive Wikimedia image-search queries from slide titles when automatic image fetching is enabled; the title/query is therefore sent to Wikimedia before a candidate image is fetched.
- **Package and tool downloads:** when you approve a dependency install, enable a component whose runner must be fetched, or approve Python/CSV analysis whose allow-listed packages are not cached, the relevant package names and versions are sent to the official registry or download host identified in the approval and Threat Model (for example npm, PyPI, crates.io, Maven Central, NuGet, or the configured component runner's source). Package contents and dependencies supplied by those services may then execute with the authority stated in the approval.
- **Google Drive integrations:** data handled by a user-installed Google Drive MCP server or integration is sent to Google according to that integration's behaviour.

- **Update check:** an automatic (on by default, opt-out) and manual check that contacts an update host we operate, served via Cloudflare, for a newer version. This is our own delivery infrastructure rather than a third-party service you configure — see section 6.4 for what it sends.
- **Cloud key check:** when an Ollama Cloud key is configured, Nyx AI makes a minimal check to Ollama Cloud (or its configured base URL) at app startup, and again whenever you save or change your Ollama Cloud settings, to verify the key still works; the save-time check also lists the provider's available models first. Where a cloud key is saved, Nyx AI may also request that provider's model catalogue (a `/v1/models` or `/api/tags` list call) to populate the model picker. When you click "Test connection" for OpenAI or Anthropic, including during onboarding, Nyx AI requests the provider's model list and makes a minimal chat request with a one-word throwaway message and a one-token output limit. If the provider declines to serve the model chosen for that test, or does not answer in time, Nyx AI repeats the same minimal request with another model from the list the provider just returned, so a working key is not reported as invalid because of which model was tested; these attempts share a single time budget of about 24 seconds and go to the same provider and key. These OpenAI/Anthropic checks occur only after that user action, not automatically at startup. None of these checks sends project content or your prompts.
- **First-use and managed browser downloads:** Nyx AI may prewarm or download a language-server binary from a documented HTTPS artifact host when you open a compatible trusted workspace and the server is enabled but not cached. Separately, if no usable installed or cached Chrome/Edge/Chromium candidate can launch, or you explicitly request setup in Settings, Nyx AI may offer to download Puppeteer's tested Chrome Headless Shell revision from Google's Chrome-for-Testing HTTPS service. A main-owned native prompt is required before the first fetch for that revision; the saved approval is bound to that exact revision, so a different tested build asks again. Every browser-archive redirect must remain the same exact revision/platform object on Google's HTTPS origin. The cached engine does not auto-update. Its first fetch is trust on first use: HTTPS protects transport, but there is no NYX-pinned upstream checksum, and a best-effort local SHA-256 can detect a later mismatch only while its local record remains intact and readable. The engine may load/render requested content and execute page scripts when used. Google/upstream terms and embedded notices apply. Managed language-server downloads follow their own availability checks.
- **MCP servers, plugins, and hooks:** plugin command files are model-visible Markdown loaded by Nyx AI and do not execute native plugin code. Stdio MCP servers and hooks — including ones defined by a plugin — run as child processes with your user permissions outside containment by default unless you enable the matching sandbox options. HTTP/SSE MCP servers are remote network services and are outside those process-sandbox controls. Any of these components may process data according to its own behaviour and trust model. An enabled MCP server can connect or reconnect automatically when its workspace is restored, and its own bootstrap command may download dependencies according to that component's configuration. Removing an MCP server also removes its stored OAuth/config credentials and pending authorisation state. A small local

removal-intent record can remain only while the two encrypted configuration/credential stores are being reconciled; the same server id is blocked until that cleanup succeeds.

[THREAT-MODEL.md](#) provides additional technical detail: it lists known network calls, destinations, data sent, and triggers. If Nyx AI's processing changes materially, this Privacy Policy will be updated before or when that change is introduced.

8. Cookies and Similar Technologies

The Nyx AI desktop application does not use cookies for tracking or analytics.

The embedded Chromium preview renders your projects. Your own project code may set cookies or local storage inside that preview. Browser-rendered web search may also receive provider cookies in its ephemeral managed profile; Nyx AI does not accept a provider's cookie banner on your behalf. Those surfaces do not create Nyx AI analytics or tracking.

Our **website (nyxai.uk)** sets no analytics, advertising, preference, language, theme, or consent-choice cookies and does not load Google Analytics. Cloudflare and Turnstile may use necessary, short-lived security technologies to deliver and protect the site and its optional forms. We will update this section before introducing any non-essential cookie or similar technology.

9. Sharing and Recipients

We do not sell personal data.

We may share correspondence data, and the website, update-delivery, and Store-report data described in this policy, with:

- **Our email provider:** the support mailbox is currently hosted at Gmail/Google. Email metadata, message content, and attachments you send may be processed by Google to provide the mailbox service and under Google's applicable terms and privacy documentation.
- **Professional advisers:** lawyers, accountants, insurers, auditors, or similar advisers where needed.
- **Authorities or courts:** where required by law or necessary to protect rights, safety, or security.
- **Successors:** if the project, assets, or rights are transferred as part of a merger, acquisition, restructuring, or similar transaction.
- **Our update host / CDN provider:** when Nyx AI checks for or downloads an update, the request reaches an update host we operate using **Cloudflare** (R2 object storage and CDN). Cloudflare processes that request — including the originating IP address and standard request metadata — as our hosting/infrastructure provider under its own terms; we use it only to deliver update files and the open-source licence-source archives described in our open-source compliance notice, not to profile or identify users.
- **Our website infrastructure and abuse-prevention providers:** Cloudflare hosts and protects the Website and provides Turnstile where forms are offered. Where configured, Upstash

Redis processes the short-lived rate-limit and replay-prevention keys described in Section 6.5.

- **Our waitlist provider:** where you join the Pro launch waitlist, **Brevo (Sendinblue SAS, France)** processes the email address and double-opt-in record for us under our service arrangement. We use the list only for the one-off launch notification.
- **Our payment provider:** where a one-time support payment is offered, **Stripe** processes the payment. Stripe acts as our processor for the payment service and as an independent controller for its own fraud-prevention and regulatory purposes under its terms. Full card numbers and CVCs do not reach our servers.

We maintain an internal supplier register recording each provider's role, the data involved, location, and transfer safeguard. In summary:

Provider	Data	Role we document	Notes
Google Workspace / Gmail	Support, legal and security correspondence	Processor or independent controller, depending on Google's contract and account configuration	Our support mailbox is provided by Google; we document Google's role and transfer safeguards in our supplier register.
Cloudflare (R2 / CDN / website)	IP address, request timestamp, app version, update-request and website access logs	Processor / service provider for our update hosting and website hosting/CDN	Update requests and website traffic are served through Cloudflare under our documented hosting/CDN arrangements.
Brevo	Waitlist email address and double-opt-in record, where the waitlist is offered	Processor for the one-off launch list	The address is added only after confirmation and is not combined with app telemetry or product accounts.
Stripe	Billing/contact details, transaction record and limited card metadata, where a support payment is offered	Processor for payment handling; independent controller for its own regulatory and fraud-prevention processing	Full card number and CVC are entered with Stripe and never reach our servers.
Upstash	Short-lived rate-limit and replay-prevention keys, where the shared store is configured	Processor for website abuse prevention	Retention and fallback limits are described in Sections 6.5 and 11.

Provider	Data	Role we document	Notes
Microsoft Store	Store acquisition, usage, health and review analytics	Microsoft is the independent controller / platform operator; we receive mostly aggregate reports, while Health may include an event-level stack trace or crash dump	Microsoft processes Store and Windows diagnostic data under Microsoft's own terms; we receive only the reports Microsoft makes available to developers.
AI providers you select	Prompts, context, code, images, embeddings	Engaged directly by you; not our processor for your project data	You engage these providers directly when you configure or use them.

AI model providers you configure. When you choose a cloud model, the prompts, context, code, and any images the model receives are transmitted **directly from your machine to that provider**. Images reach the model in two ways: attachments you supply or ask it to analyse, and screenshots Nyx AI captures automatically during preview or build work so the model can be told what the page looks like. The provider request may also include Nyx AI's system instructions and relevant tool definitions — these describe how the model should respond and which tools are available. If your selected model cannot process images, Nyx AI may hand the image to the vision provider and model resolved by the application's vision settings (shown in the app) so it can return a description — unless the 'Keep local-model images on this device' setting blocks the local-to-cloud case; if the provider or an unambiguous destination cannot be resolved, the image remains on your device. Depending on the model you select, that recipient is:

- **Anthropic** or a **custom Anthropic-compatible endpoint** you configure (Claude-compatible models),
- **OpenAI** or another **custom OpenAI-compatible endpoint** you configure (GPT / OpenAI-compatible models),
- **Ollama Cloud** (ollama.com) or a custom Ollama Cloud base URL, when you use a cloud model or cloud-backed semantic index,
- or a user-configured **remote Ollama endpoint**, including an endpoint labelled Ollama Local that is not loopback.

Each provider processes that data under its own terms and privacy policy and may act as an independent controller, your processor, or another role under its own terms. Primary model selection is provider-qualified. Data goes only to the selected provider/endpoint, a visible same-provider model fallback when the exact model becomes unavailable during a run, or the vision handoff resolved and permitted by the application's vision settings; Nyx AI does not silently switch the request to a different provider. If you run a **local** model and no remote same-provider route or cloud vision handoff is used, nothing in this category leaves your machine.

Third-party AI providers and integrations you configure are not our subprocessors for your project data. You engage them directly by configuring or using them.

10. International Transfers

Correspondence sent to the Gmail support address may be processed by Google in the UK, EEA, United States, or other locations as described in Google's documentation and transfer safeguards.

Update-check and update-download requests, visits to our website, and Turnstile form-protection signals are served through Cloudflare, which is headquartered in the United States and operates a global edge network, so a request (including its IP address) may be processed in the United States or other locations under Cloudflare's terms and transfer safeguards. Waitlist data held by Brevo is stored within the European Union; transfers from the UK to the EEA are permitted under the UK's adequacy regulations. Stripe is headquartered in the United States and may process payment-related personal data there and in other locations. Upstash may process the short-lived abuse-prevention keys on the regional or other infrastructure documented for our configured service. Where those transfers are ours to safeguard, the applicable mechanism is recorded in our supplier and transfer register.

For each current and future processor, we record in our internal supplier and transfer register which transfer mechanism applies — a UK adequacy route or "data bridge", the UK International Data Transfer Agreement (IDTA) or the Addendum to the EU Standard Contractual Clauses, or another lawful safeguard — together with a transfer risk assessment where we rely on safeguards. You may request a copy of the relevant safeguard for a transfer we make (with commercially sensitive information redacted where necessary) by emailing customerservice@nyxlimited.com.

For data you send directly to third-party providers through Nyx AI, international transfers depend on the provider, endpoint, account, settings, and content you choose. You are responsible for assessing those transfers for your own projects and users.

If we later appoint processors for hosted services, support tooling, additional payment services, telemetry, or accounts, we will update this policy with the relevant categories and safeguards, such as UK adequacy regulations, EU adequacy decisions, the UK International Data Transfer Agreement/Addendum, or EU Standard Contractual Clauses where appropriate.

The UK currently has EU adequacy for transfers from the EEA to the UK (renewed on 19 December 2025 and valid until 27 December 2031). That does not remove the need to assess other transfers made by third-party providers you choose.

11. Retention

We keep personal data only for as long as needed for the purposes described in this policy.

Data category	Retention
Support, legal, security, accessibility, and other correspondence	Up to 2 years from the last meaningful contact, unless a longer period is needed for security, dispute, legal, or abuse-prevention reasons
Data rights requests and data-protection complaints	Up to 6 years, to evidence handling of the request or complaint
Update checks (automatic and manual)	No application-level record kept by us; the request reaches an update host we operate via Cloudflare and appears only in Cloudflare's standard server/CDN access logs, which NYX LIMITED does not separately store and which Cloudflare retains under its own documented retention
Website access/security request metadata (Cloudflare)	Not separately stored by us; retained by Cloudflare according to the service configuration and agreement that apply to our website hosting
Pro waitlist email and double-opt-in record, where offered	Until the launch notification is sent and for up to 3 months afterwards, or until you unsubscribe, whichever is sooner. Unconfirmed addresses are not added to the contact list; Brevo confirmation links expire after 30 days, while the submission attempt and confirmation email may remain in Brevo's transactional logs under its documented retention
Website rate-limit and replay-prevention keys, where Upstash is configured	Short-lived and automatically expired no later than 24 hours after the relevant protection window
Website support-payment and transaction records, where offered	Up to 6 years after the transaction to meet tax, accounting, and dispute-handling obligations
Microsoft Store developer reports (Store installs; Health reports can include an event-level stack trace or crash-dump file)	Event-level Health data we download is retained only while the referenced defect or reliability issue is being investigated, then deleted; the reports themselves remain available in Microsoft's developer tools under Microsoft's own retention
Local app data	Stored on your machine until you delete it. Uninstalling Nyx AI does not delete it.

For transparency about what is written where: in Nyx AI's **chat mode**, model tools cannot run shell commands or directly read or write arbitrary workspace paths. Generated previews and downloadable artifacts are carried in the chat and may therefore be retained with chat history in Nyx AI's local user-data folder. Document rendering may use short-lived temporary files that Nyx AI removes on a best-effort basis. Scheduled Chat tasks write durable output to `Documents\Nyxion\Scheduled` and may add downloadable artifact payloads to chat history; these

writes occur under the schedule you created, without a fresh click when it fires. Saving an ordinary artifact to your workspace or opening it externally otherwise writes only after your explicit action. Nyx AI also keeps local chat history, optional long-term memories, settings, audit data where enabled, a small usage counter, a small window-layout file (`window-state.json` , the app window's last size and position), and local crash and non-fatal-error records (version-stamped, secret-scrubbed before writing, rotated and capped, never transmitted, and removed by "Delete all my data"); none of that product data is sent to Nyx Limited automatically (see section 4.2). In **code mode**, the assistant edits files in the workspace you open, which is the purpose of that mode.

On Windows machines where Locked execution is available, Nyx AI also keeps an encrypted local grant-recovery journal in its user-data folder. It can contain one-way workspace identifiers, Nyx AI-created sandbox identifiers, timestamps, and the exact local workspace, allowed-folder, interpreter, scratch, or temporary paths to which sandbox access was granted. Its encryption key is protected through the operating system's secure storage; there is no plaintext fallback. Nyx AI uses this state only to reconcile and revoke its own access grants after a crash, settings change, workspace switch, or data deletion. Closing Nyx AI does not itself remove those grants: quitting records them for verified removal at the next launch, so a grant can persist on disk between sessions. The journal is never transmitted to NYX LIMITED or any third party. Verified cleanup removes the corresponding record; unresolved cleanup details and quarantined corrupt journal bytes may remain locally so the app can retry or report them honestly until "Delete all my data" removes the journal, its protected key, and its temporary/quarantine files. If the journal cannot be securely read or persisted, Nyx AI does not create a new Locked-execution access grant under that identity.

You can delete local Nyx AI data from your machine. Deleting local data does not delete correspondence you previously sent to us by email, but you may contact us to request deletion where applicable.

The in-app "Delete all my data" action removes the known, enumerated set of files and folders Nyx AI itself creates in its user-data directory (including chats and their recovery copies, memories and vector indexes, encrypted settings recovery copies, the encrypted AppContainer grant-recovery journal and protected key, browser-engine storage cleared through the engine's own clear-storage API, and Nyx AI's temporary files); it works from that enumerated list rather than indiscriminately emptying the folder, and the enumeration is maintained to cover everything Nyx AI writes there. Before permission cleanup, Nyx AI stops running Locked-execution commands and quiesces sandboxed integrations so no contained child retains the authority being revoked. If safe cleanup cannot be proved after a child or integration was stopped, local app data is not deleted and that integration may deliberately remain disconnected; reconnect it or restart Nyx AI before retrying. It also removes attributable Nyx AI AppContainer profiles; the app-managed language-server cache at `.nyx/lsp-servers` under your Windows profile; each recorded workspace's exact `.nyx/` metadata and `.nyxdeps/` dependency/cache directory; and the whole dedicated `Documents\Nyxion\Scheduled` folder used by scheduled Chat tasks. Crash-stale Nyx AI GUI-desktop grant markers are reconciled and their corresponding desktop access-control entry is revoked before the records are erased. Nyx AI keeps a small local registry containing opened workspace paths and the first and most recent

times each was recorded, solely so the delete-all action can find these app-managed workspace directories; the registry itself is deleted by the same action. Workspaces last opened only on app versions before this registry existed may not be in it and are then not reached — reopen them once (or manually delete only their `.nyx/` and `.nyxdeps/` directories) to include them. If the current registry is unreadable or corrupt, deletion reports that it cannot prove coverage instead of silently claiming success. Outside the exact app-managed folders just listed, the action does not delete project source. Because those listed directories are removed as whole app-data locations, any files you manually place inside `.nyx/`, `.nyxdeps/`, `.nyx/lsp-servers`, or `Documents\Nyxion\Scheduled` are removed too. Anything that cannot be removed in the moment (for example a file locked by another program, an offline grant target whose revocation cannot yet be verified, or engine state files the browser engine recreates empty on exit) is never silently counted as deleted: the completion dialog tells you that leftovers remain and lists the first of them (it shows up to 12 lines and 2,500 characters of errors and leftover items combined). That leftover list covers the enumerated locations the action attempted: a file inside the user-data folder that Nyx AI does not recognise as its own (for example one another program placed there) is not deleted and does not appear in the list.

The assisted uninstaller deliberately preserves local app data so that uninstalling or reinstalling does not unexpectedly destroy chats, settings, memories, caches, or workspace metadata. **Uninstalling is therefore not the same as deleting your Nyx AI data.** The same is true of the filesystem permissions Nyx AI grants to its own sandbox identity so that Locked execution can reach your workspace, allowed folders, interpreters and caches: those access-control entries are tracked in the journal described above and are removed only by a confirmed "Delete all my data", so they survive an uninstall. If you want that data and those permissions erased, use "Delete all my data" before uninstalling. The same main-owned deletion action is available from the legal-consent screen before you accept updated terms; Windows asks for a separate native confirmation before deletion begins. If the app cannot be opened, remove the app's user-data folder and the exact app-managed locations described above manually — including the GUI-desktop grant-marker folder at

`%LOCALAPPDATA%\Nyxion\gui-desktop-grants`, any `nyx-render-*.html` temporary render files in your temp folder or in a `nyxion\render-tmp` folder at a drive root, and the sandbox identity folders named `Nyxion.Sandbox.*` under `%LOCALAPPDATA%\Packages` together with their matching entries under the `AppContainer` mappings in your user registry hive — taking care not to delete project source.

What "attributable" excludes, stated plainly. "Delete all my data" removes the sandbox identities it can attribute to a workspace through its own permission journal. Identities created before that journal existed cannot be attributed to anything, and are deliberately **not** inferred to be safe to remove, because deleting one whose filesystem permissions had not first been proved removed would leave those permissions behind naming an identity that no longer resolves. Those are retained and reported to you by name in the deletion result rather than silently kept. They contain no chat, file or credential content — a folder name carries a one-way truncated hash of a workspace path, and the folders themselves are a few kilobytes — but they do persist after deletion, and on a machine used heavily with Locked execution there can be well over a hundred of them. Remove them by hand using the locations above if you want them gone.

12. Security

Nyx AI is designed to keep project data local by default. Local safeguards include operating-system secure storage for supported secrets, approval flows for risky actions, local audit logging where enabled, and documented network boundaries. Two separate local logs exist and they are not equally protected: one is encrypted and hash-chained so edits to it are detectable, while the separate activity log is an ordinary plaintext file on your machine that records the text of commands the assistant issued. It is secret-scrubbed on a best-effort, pattern-based basis, it is not encrypted and not tamper-evident, and its rotated archives are not deleted automatically. Both are local and neither is transmitted; the difference is set out row by row in `THREAT-MODEL.md`.

For correspondence we receive, we use reasonable organisational measures, including limiting mailbox access and using the security controls provided by the email service.

No method of electronic storage or transmission is perfectly secure. Please do not email secrets, API keys, private repository contents, or highly sensitive personal data unless strictly necessary. If a vulnerability report needs sensitive material, send only the minimum needed to reproduce the issue.

If we become aware of a personal data breach that is likely to result in a risk to your rights and freedoms, we will notify the ICO within 72 hours where required. If the breach is likely to result in a high risk to your rights and freedoms, we will also notify affected individuals without undue delay where required.

13. Your Rights

Depending on the circumstances and lawful basis, you may have the following rights:

- The right to be informed
- The right of access
- The right to rectification
- The right to erasure
- The right to restrict processing
- The right to data portability
- The right to object
- Rights relating to solely automated decision-making with legal or similarly significant effects
- The right to withdraw consent where processing is based on consent
- The right to complain to a supervisory authority

To exercise your rights, email customerservice@nyxlimited.com with the subject line "Data Rights Request".

We may need to verify your identity. We will respond without undue delay and within one month unless the law allows an extension. If your request is complex or you make multiple requests, we may extend by up to two further months and will tell you within the first month. If we reasonably need information to confirm your identity, the response period begins when we receive it. Where the law

allows us to seek clarification of your request, the response period is paused only while we await your reply; we will ask promptly and will not request more information than is necessary.

The right to object must be brought to your attention clearly and separately: **you may object to processing based on legitimate interests at any time by emailing customerservice@nyxlimited.com**. We will stop the processing unless we have compelling legitimate grounds or need the data for legal claims.

14. Data-Protection Complaints

We operate a formal **Data Protection Complaints Procedure** with a structured complaint form and several routes. To raise a complaint, use that form (the procedure and form are included in the application's Legal panel and linked here), or email customerservice@nyxlimited.com with the subject line "Data Complaint", or write to us at the registered office above.

This implements the Data (Use and Access) Act 2025 complaints-handling requirement (in force **19 June 2026**): we provide a clear, accessible way to complain, **acknowledge your complaint within 30 days**, investigate without undue delay, keep you informed of progress, and tell you the outcome without undue delay; and we keep records of complaints.

You can complain to the ICO at any time:

- Website: <https://ico.org.uk>
- Phone: 0303 123 1113

If you are in the EEA, you may also complain to your local supervisory authority.

15. Children

Nyx AI is a developer tool for people aged 18 or older and is not offered to children. We do not knowingly collect personal data from children through the application.

If you believe a child has sent us personal data, contact customerservice@nyxlimited.com and we will take appropriate steps.

16. Automated Decision-Making

We do not use personal data to make solely automated decisions about you that produce legal or similarly significant effects.

AI models you connect to through Nyx AI generate code, text, or other output at your direction. They do not make decisions about service access, pricing, support, or rights on our behalf.

17. Data Protection Fee and Registration

UK organisations, sole traders, and other controllers that process personal data may need to pay the ICO data protection fee unless an exemption applies. Whether a fee is due depends on the controller,

processing purposes, and exemptions.

Current status: NYX LIMITED is registered with the Information Commissioner's Office (ICO). ICO registration number: ZC167877.

Even if an exemption applies, data protection obligations still apply.

18. UK / EEA Representative

No UK representative is required where the controller is established in the UK.

Whether NYX LIMITED must appoint an EEA representative depends on whether particular processing falls within Article 3(2) EU GDPR — which can apply where goods or services are offered to people in the EU, including potentially free services, depending on evidence of targeting — and whether the limited Article 27 exemption applies. The desktop app operates no accounts, telemetry, advertising, or behavioural monitoring, and the UK-operated website uses no analytics; it may offer the optional waitlist and support-payment flows described above. We review this assessment when Nyx AI's EEA distribution, marketing, website activity, integrations, or commercial services change.

19. California Residents

We do not sell personal information or share personal information for cross-context behavioural advertising.

In the preceding 12 months, the categories of personal information we may have collected, as described elsewhere in this policy, are:

- **Identifiers and correspondence content** — name, email address, and message content if you email us, plus a waitlist email address if you choose to join where it is offered ([Sections 4.1 and 4.4](#)).
- **Internet or other electronic network activity information** — the IP address and app version seen by the update host we operate for update checks ([Section 6.4](#)), and website request/security, Turnstile, and short-lived rate-limit information processed through Cloudflare and Upstash where configured ([Section 6.5](#)).
- **Commercial information** — billing name and email address, transaction record, and limited card metadata received from Stripe if you make an optional support payment where offered.
- **Reliability diagnostics received from Microsoft** — for Microsoft Store installs, developer reports that are mostly aggregate but whose Health data can include an event-level stack trace or crash-dump file ([Section 6.1](#)).

The sources, purposes, sharing, and retention for each category are described in Sections 4 to 6, 9, and 11. To exercise California privacy rights, email customerservice@nyxlimited.com with the subject line "CCPA Request".

This section is included because software may be downloaded internationally. It does not mean that every California privacy law applies to every release or user.

20. Changes to This Policy

We may update this policy to reflect changes in Nyx AI, our processing, or the law. Material changes should be notified in-app, in release notes, or through the distribution channel.

The policy supplied with a particular version describes the processing for that version unless a later policy clearly applies.

21. Contact

For privacy questions, rights requests, legal notices, complaints, security reports, or accessibility issues:

Email: customerservice@nyxlimited.com

22. Governing Language

Subject to applicable mandatory law, the English-language version of this Privacy Policy governs. Any translation of this policy, of the [Terms](#), the [LICENSE](#), or of the Nyx AI application interface is provided for convenience only, and in the event of any conflict, ambiguity, or inconsistency the English-language version prevails. Where mandatory law gives you rights based on information provided in another language, nothing in this paragraph limits those rights.

Nyx AI - Code with intelligence. Nyx AI transmits locally stored application data only through its documented network features — the default-enabled update check, third-party services you configure, trigger, or enable (such as a cloud model provider, web fetch/search, marketplace fetch, Google Drive integration, first-use tool download, or image/vision handoff), and the automatic requests those enabled features make. Those flows go to the relevant provider, documented infrastructure host, or to us only when you choose to send correspondence or a report. Separate processing by Windows, the Microsoft Store, your project code, or other software is governed independently, as explained in the Privacy Policy.