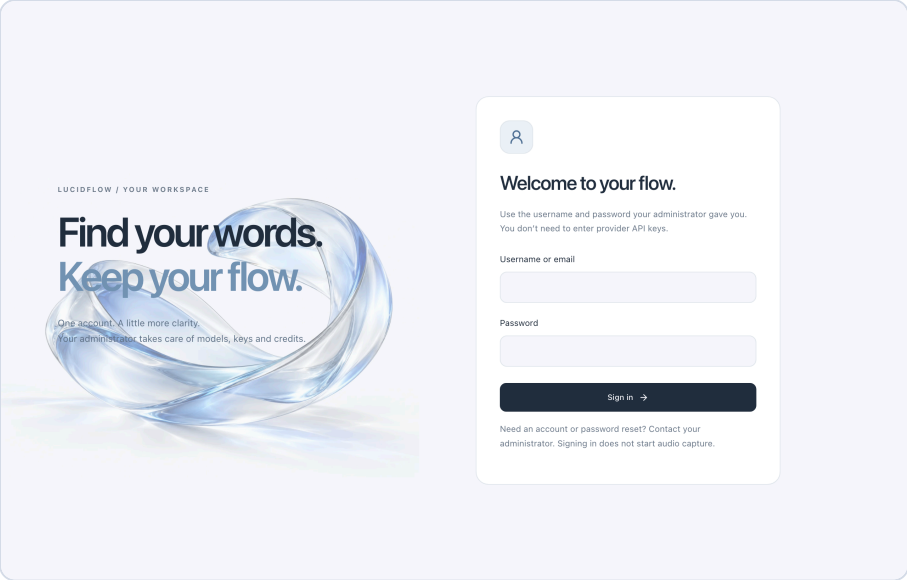
LUCIDFLOW DESKTOP · V1.7.1 · 6 NOVEMBER 2026

Find your words. Keep your flow.

The click-by-click guide to your cloud-managed desktop.



The current app opens as Lucidflow. Example test account screens are used in this guide.

Your first five clicks

1. Install the current Windows EXE or Apple Silicon Mac DMG.
2. Read consent, then click **Agree and continue**.
3. Sign in with the username and password your administrator gives you.
4. Open **Tutorial**, then add résumé and job context in **Setup**.
5. Check your account access/credits and test audio before **Live Session** → **Start session**.

You do not enter provider API keys in the managed app. Your administrator creates your account, assigns keys, adds credits and manages allowed models and shortcuts.

Guide: installation/account p2 · setup/models p3 · live audio p4 · popup/keys p5 · practice/history p6 · administrator p7–8 · troubleshooting/privacy p9.

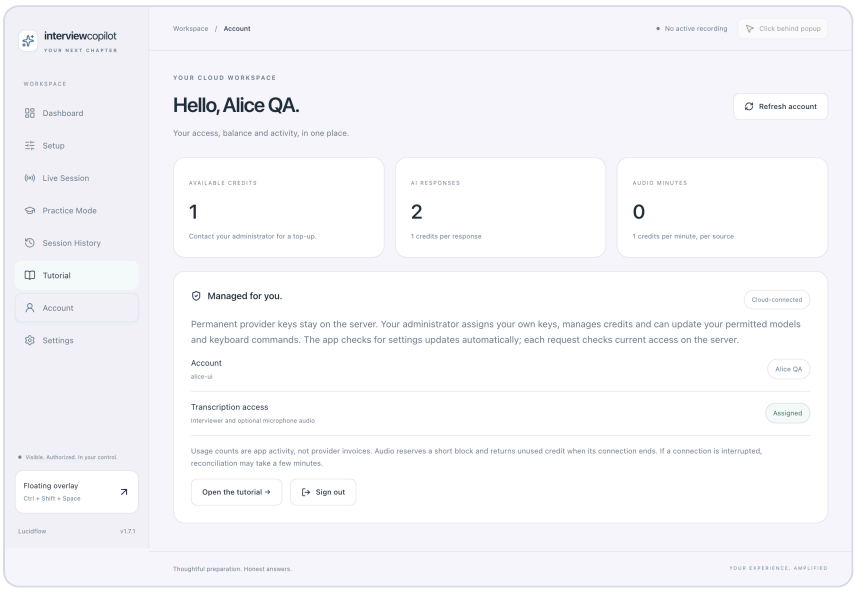
Install. Sign in. Find your workspace.

Windows x64

1. Open **Lucidflow-Setup-1.7.1-x64.exe**.
2. Choose an installation folder and click **Install**.
3. Open the desktop shortcut, accept consent and sign in.

Apple Silicon Mac

1. Quit the old app. Open **Lucidflow-1.7.1-mac-arm64.dmg**.
2. Drag Lucidflow into **Applications**.
3. Open it from Finder, accept consent and sign in.

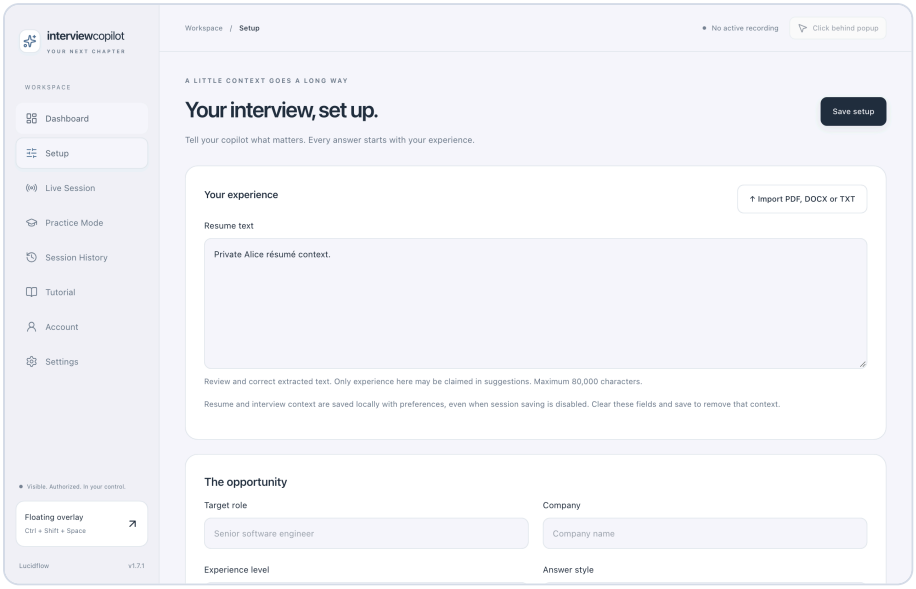


The Account screen shows app credits and usage. It does not reveal permanent provider keys.

- **Refresh account** checks balance, access and current cloud settings immediately. Normal synchronization also runs automatically.
- **Sign out** stops the current session, hides the popup and removes the local login. A different cloud user gets separate context/history.
- Need a login, password reset, keys or credits? Contact your administrator.

Both installers are unsigned; the Mac build is not notarized. Check the checksum and release details before installing. Mac computer audio requires macOS 14.2+ and recording permission; Windows runtime needs real-device acceptance testing.

A little context makes it personal.



Your résumé and job context stay in account-specific local preferences until used in a provider request.

1. Open **Setup**. Click **Import PDF, DOCX or TXT**, or paste into **Resume text**.
2. Fill **Target role**, **Company**, **Experience level**, **Job description** and optional coaching instructions.
3. Choose answer style, target answer length, transcription language and the audio sources you intend to use.
4. Click **Save setup**. Stop the session before changing interview/audio settings.

Choose a permitted AI model

Use **Choose your AI model** on Dashboard, Live Session or Practice Mode. Provider/model changes save immediately. Only assigned providers and allowed models can be used. A locked model is controlled by your administrator. The managed Setup screen directs model changes to this picker.

> **No keys to paste.**
>
> Open **Account** if access is pending. Your administrator assigns a separate AI key and transcription key to your account in the cloud.

Let the interviewer’s question come to you.

1. Join the call in Teams, Meet, Zoom or another app on this computer. Keep its audio playing locally; headphones help separate your voice.
2. Open **Live Session**. Choose **Computer audio · remote interview call** for interviewer audio. An interviewer input device can be used for a configured loopback source.
3. Open **Setup → Audio diagnostics → Test interviewer audio**. Play the call audio, then listen to the local sample. This test does not send audio to providers.
4. On Mac, allow **System Audio Recording** when requested. A microphone source is separate and represents your own voice.
5. Click **Start session**. Recording status becomes visible. The floating popup opens after capture starts.
6. Wait for speech. Interim transcription appears, the complete question is detected, and a suggestion streams into the session and popup. You do not need to type it.
7. Click **Stop session** when finished. Hiding the popup or stopping answer generation does not stop recording.

When you need a manual question

Expand **Manual question (optional)**, enter text and click **Answer typed question**. This is optional; live question detection remains the normal flow.

Control	What it does
Answer now / Regenerate	Generate or replace the current suggestion; each makes an AI request.
Shorter / Longer	Request a different answer length; uses AI credits.
Stop generation / Clear	Cancel the current answer, or clear it. Neither stops capture.
Reconnect transcription	Reconnect after a transcription problem, once access/network/credit is fixed.

Actual meeting-service behavior, physical audio and permissions must be checked on your device. The app does not join the meeting as a participant.

Keep the popup. Click underneath.

1. Use **Floating overlay** in the sidebar, **Open overlay**, or start a live session.
2. Drag the popup's title bar to move it; drag an edge/corner to resize it while it is interactive.
3. Click **Click behind popup** in its header or the main top bar. Clicks and scrolling now reach the window underneath while the popup remains visible.
4. Restore with **Interact with popup** in the main app or your click-through shortcut. Hide/show also resets this mode.

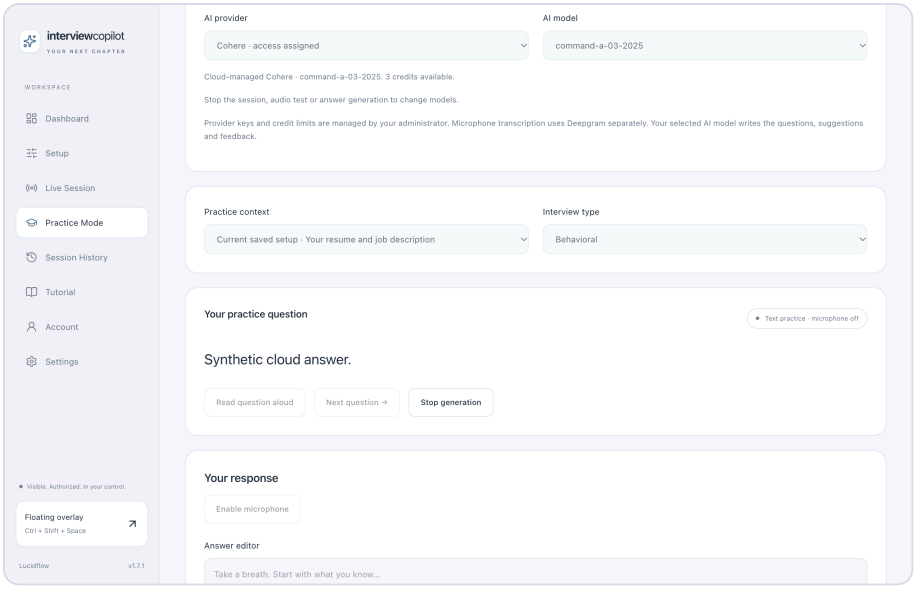
All seven default commands

Press	Action	Result
Control + Shift + Space	Show / hide popup	Hide resets click-through; capture can continue.
Control + Shift + A	Answer now	Use the current interviewer question.
Control + Shift + R	Regenerate	Request a fresh answer.
Control + Shift + J	Shorter answer	Request a concise version.
Control + Shift + K	Longer answer	Request a more detailed version.
Control + Shift + X	Stop generation	Stop answer generation, not recording.
Control + Shift + C	Click-through	Click/scroll underneath; press again to interact.

> **Your tutorial shows the actual keys.**
>
> Open **Tutorial → Every shortcut**. The administrator can change these bindings in the cloud, so your actual keys may differ. Defaults use **Control** on Mac too, not Command. Conflicting OS shortcuts keep the previous working bindings.

Answer, regenerate, shorter and longer each use AI-request credits. A shortcut never starts recording remotely.

A little rehearsal. A little perspective.



Example question from a synthetic test provider. Your account uses its permitted provider.

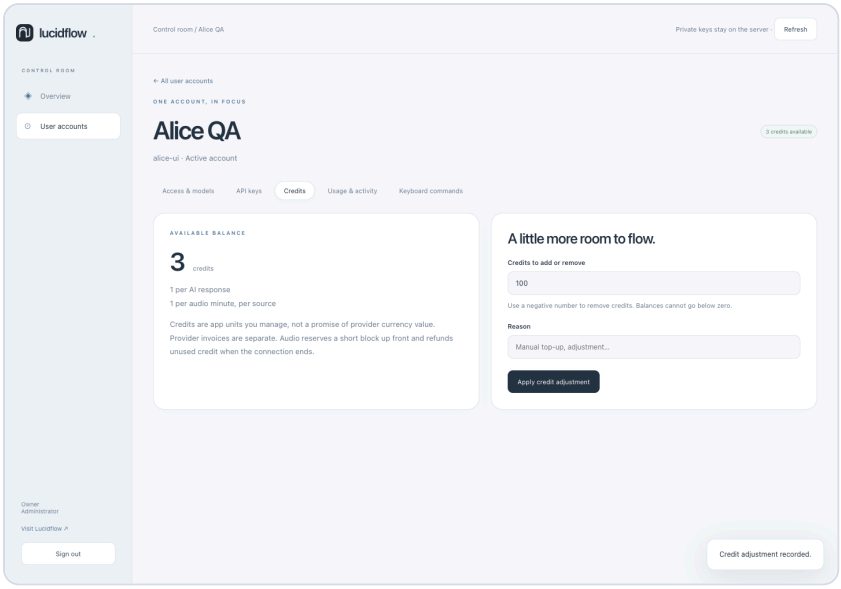
1. Open **Practice Mode**. Choose interview type and current/saved context.
2. Click **Start text practice**, or **Start microphone practice** if you want voice transcription and have audio access.
3. Read the question. **Read question aloud** uses a local speech voice when available. **Next question** makes another AI request.
4. Type into **Answer editor**, or use your microphone transcript. Click **Get feedback**.
5. Click **End practice** when finished.

Review what you keep

In **Session History**, use Search history, select a session, then **Export Markdown**, **Export JSON** or **Delete session**. Delete-all requires confirmation. Session saving is optional; turn it off in Settings for new sessions. Previously saved sessions remain until deleted.

Each cloud account uses its own local history/context. The cloud usage log stores counts and credit changes, not the transcript or answer contents.

For the owner: one calm control room.



Owner-only dashboard. The example account is test data; no permanent keys are displayed.

1. Open **https://lucidflow-xi.vercel.app/admin/**. First owner setup uses the private page opened by `./scripts/pnpmw cloud:admin-setup` from the project with your Vercel owner login. Choose your admin username/password there.
2. Click **Create user**. Enter their display name, username/email and initial password (12+ characters). Share the login securely.
3. In that user's **API keys**, assign their own provider key and Deepgram key. Saved keys are not returned to the UI.
4. In **Credits**, enter a top-up and reason; click **Apply credit adjustment**. Users do not buy credits online in this release.
5. In **Access & models**, set defaults, permitted pricing, output budget, enabled status and optional model lock.

Create a separate user account for your own desktop use. The administrator account is for managing users. Do not paste provider keys into the desktop or share the admin login with customers.

Rotate access. See usage. Set every command.

Admin section	What to click
API keys	Enter a key and click **Assign key**. Enter a replacement and click **Rotate key**. **Remove** revokes the assignment.
Allowed model IDs	Use one model per line or commas; click **Save allowed models**. Keep the default in its allowed list.
Access & models	Set per-request and per-audio-minute app credits, default provider/model and output budget; click **Save access & pricing**.
Account enabled	Disable and save to revoke sessions. Reset password also signs the user out on all devices.
Keyboard commands	Edit any of seven bindings and click **Save keyboard commands**, or **Use defaults**. Desktop checks approximately every 20 seconds.
Usage & activity	See latest 100 requests, key versions, character counts, audio duration, credit ledger and admin actions.

How app credits work

- Default pricing is **1 credit per AI request** and **1 credit per audio minute, per source**. The owner can change it per user. Credits are app units, not provider currency/token invoices.
- Questions, suggestions, regeneration, feedback and optional AI classification make AI requests. A request that delivers no answer is refunded; partial delivered answers consume the configured request cost.
- Audio reserves 30-second blocks per source before forwarding PCM. The balance must cover the initial reservation. Unused reservation returns when the connection ends; interrupted server leases reconcile within a few minutes.
- Both key assignments and account access are checked by the server. New requests use the current key without rebuilding or restarting the EXE/DMG. Active audio periodically rechecks access and reconnects for a rotated key.

Permanent keys are encrypted per user/provider. User balances cannot go below zero, and repeated adjustment/request IDs do not charge twice. Provider account credit, model access and invoices still need to be managed at the provider.

If something interrupts your flow.

What you see	What to do
Cannot sign in	Check the username/password with your administrator. A disabled account or password reset revokes sessions.
Access pending / no credits	Ask the administrator to assign the required provider keys or top up credits, then click **Account → Refresh account**.
No interviewer question	Choose call/computer audio, play call sound locally, test interviewer audio in Setup, and check recording/transcription status. Your microphone is not the interviewer source.
Transcription key/access error	Ask the administrator to correct or rotate the audio key; refresh your account, then reconnect transcription or stop/start the session.
Shortcut conflict	The old working bindings remain. Ask the owner to change conflicting keys; see the actual active keys in Tutorial.
Cannot click popup	Use **Interact with popup** in the main app, your click-through shortcut, or hide/show.
Startup/build error	Quit the old process and reopen the current installer. In development rebuild main/preload/renderer with `./scripts/pnpmw build`, then restart.

Privacy and recording

Capture begins only after your action. Audio and prompts pass through the cloud service to external providers. The app does not save raw call audio. Optional local history contains transcripts/answers; the cloud keeps account/usage/credit metadata, encrypted keys and password hashes. Your device encrypts its sign-in session using OS-backed storage.

Screen sharing

Native window protection does not guarantee invisibility in every OS or sharing tool, especially Mac ScreenCaptureKit. Check the receiving view. Use assistance and recording only where permitted.

Validated with synthetic provider/audio responses on a Mac and a separate cloud test database. Windows physical runtime, real-provider access, actual meeting calls and device permissions still need verification on the intended setup. Existing legacy local history/context is retained separately during migration.