Cloudroom

SECURITY OVERVIEW

Version 2.3 · October 16, 2026

How we protect your code and keys.

Current controls and limitations for our desktop app and cloud service.

IN SHORT

Built for agents with real access. Isolated, encrypted, and open source.

This document describes controls we run today. We do not have a SOC 2 report yet.

Cloud isolation	**One sandbox per thread**
Stored keys	**AES-256-GCM**
Core and desktop app	**Open source**
Main servers	**United States**

THE DATA WE HANDLE

DATA	WHAT IT IS	HOW WE HANDLE IT
Account	Email and sign-in provider	Deleted within 30 days of closing your account.
Billing	Plan and status	Polar takes payment. We never see your card details.
Cloud projects	Code and **.env** files you send	Kept in the thread's sandbox and its backups. Deleted with the thread.
Chat history	Cloud thread messages	Unused threads auto-delete: 60 days on Free, 1 year on Pro, 5 years on Ultra.
Logins & API keys	Ones you copy to the cloud	Encrypted in our database. Used only to run your threads.
Usage data	App usage and setup errors	Never prompts, code, or files. Turn off in Settings.
Mac commands	Cloud agents' Mac commands	Saved in chat history. An outside AI checks them for risk, with secrets removed where possible.
Local threads	Run on your Mac	Stored on your Mac. AI titles and voice notes use the standard AI inference providers. Agent feedback and share links go to us.

CONTROLS IN PLACE TODAY

01 Isolation

- Each thread gets its own sandbox. Its subagents can share it.
- A sandbox's access token works only for that sandbox.
- Cloud cores require a secret 256-bit token and HTTPS.

02 Your models, your accounts

- Agents call Anthropic, OpenAI, and others directly, with your own logins. We never sit in between.
- We never use your content to train AI models.
- We never sell your data.

03 Encryption

- HTTPS/TLS to our servers and sandboxes.
- Saved logins and API keys are encrypted with AES-256-GCM before storage. Only our server can decrypt them.
- Row-level security blocks all browser access to them.
- Our database is encrypted at rest (AES-256).

04 Access control

- Only named admins can change the prod database.
- Our own AI agents get read-only database access.
- Secrets live in a password manager, never in code.
- Sign in with Google, GitHub, or email. Breached passwords are blocked.

MORE CONTROLS

05 Secure development

- Core and desktop app are open source (Apache 2.0).
- Multiple AI reviews are running 24/7 on the codebase, alternating between frontier models from leading AI labs.
- CI scans every commit for leaked secrets.
- Mac releases are code-signed and notarized by Apple.

06 Monitoring & recovery

- Health checks alert the team every hour.
- AI flags risky commands cloud agents run on your Mac.
- Daily database backups, kept for 7 days.
- Incidents get written root-cause reports.

YOU STAY IN CONTROL, AND THE LIMITS WE KNOW ABOUT

07 You stay in control

- Mac access has four levels: Off, Read-only, Ask first, and Full. Change it in one click.
- Command Guard, on by default, blocks catastrophic commands, like wiping a disk or deleting a GitHub repo.
- When an agent needs an API key, you enter it in a secure form. It never appears in chat, history, or logs.
- Share links are read-only snapshots you create. Anyone with the link can read it. Stop sharing deletes them.
- Choose when unused cloud threads auto-delete.
- Choose which GitHub account cloud agents use. Disconnect it anytime.
- Turn off usage data and agent feedback in Settings.

! Known limits

- Agents run without approval prompts. With admin rights in their sandbox, they can read anything there, including keys stored in it.
- Sandboxes have open internet access, so agents can install tools. We don't filter that traffic yet.
- Removing secrets from share links and command reviews is best effort.
- Not yet available: SSO/SAML, audit-log export, customer-managed keys, and an external penetration test.

SECURITY REVIEW ANSWERS

Who on our staff can access your data?
Our developers access account data, cloud chat history, and diagnostics only to fix problems or support you.

Data processing agreement (DPA)
We sign one on request.

Where data is processed
Database and website: US West (California). Cloud sandboxes: mainly the United States. File storage: Western North America.

Two-factor login
Sign in with Google or GitHub to use their two-factor login. We don't offer our own yet.

Incident notification
If an incident affects your data, we email you without undue delay.

What backups cover
Database: daily, kept 7 days. Cloud threads: sandbox snapshots and file backups, until the thread is deleted. Local threads stay on your Mac.

SHARED RESPONSIBILITY

Cloudroom secures

- Isolation between sandboxes and between users
- Your saved logins and keys, encrypted at rest
- Our servers, database, and app releases
- Incident response and fixes

You decide

- What agents can reach: Mac access level, keys, and repos you connect
- Whether Command Guard stays on
- Reviewing what agents change before you ship it
- Your model accounts, under your provider's data terms

Running a security review?

Ask us for an architecture and data-flow overview, or answers to your security questionnaire. Our service providers are listed in our privacy policy.

Security contact	**david@cloudroom.dev**
Security page	**cloudroom.dev/security**
Privacy policy	**cloudroom.dev/privacy**
Source code	**github.com/davidondrej/cloudroom**

This document describes the security practices of Cloudroom Inc. as of October 16, 2026. It is not a SOC 2 report, certification, or audit opinion. Our database provider (Supabase, supabase.com/security) and hosting provider (Vercel, security.vercel.com) hold their own SOC 2 Type II reports. Those cover their infrastructure, not Cloudroom's own controls.