

Product Policy Document: Cursor with Privacy Mode (harness)

Product and vendor: Cursor, Anysphere Inc.

Way of working: Harness. An agentic environment that works over a folder of files the firm controls, bringing the model to the work instead of uploading the work to a platform.

Status: Ready to adopt, subject to the signup-time confirmations in Section 3. Cursor is the benchmark harness in the firm's presentation materials, and the approved configuration below (Privacy Mode on, admin-enforced) is the configuration those materials describe.

Implements: Model Firm AI Use Policy (42/02). Exhibits: 07 (Cursor data-use page, retrieved 2026-08-07); 14 (Cursor security page, retrieved 2026-08-07); 12 (Anthropic Covered Models retention article, retrieved 2026-08-07); 11 (Wave 1 exhibits pass memo).

Document date: 2026-08-07.

1. What this is

Cursor is a working environment: a program on the firm's machines that opens a folder, reads the files in it, and puts an AI agent to work over them, drafting, editing, comparing, and organizing, with every change visible as a change. The documents stay in the folder. What goes to the vendor, per request, is the assembled context: the specific text the harness collects for that one task, sent over an API channel.

The confidentiality architecture is the reverse of a platform. With **Privacy Mode** enabled, Cursor states that it will not train on Customer Data and that it "maintains zero data retention (ZDR) agreements with all providers, and AI model providers will not store or train on your data" (Exhibit 07). The working set, the instructions, the session logs, and the memory all live in the matter folder on firm storage. The trade is agency for control: the agent can read more than you meant and act faster than you watch, so the policy's scoping and human-review rules do the work that a vendor's walls do on a platform.

Two caveats belong next to the ZDR claim, in the same breath. First, **the abuse-detector exception**: even in Privacy Mode, model providers (and Cursor) may run risk classifiers, and prompts that trigger abuse detectors may be stored for investigation under the provider's retention policy (Exhibit 07). Second, **the covered-model safety window**: Anthropic's newest covered models carry a 30-day safety retention even for ZDR customers; the model stays off until an administrator accepts those terms, and the window covers only requests actually routed to that model (Exhibits 07 and 12). Route work to non-covered models under ZDR and nothing is persisted.

2. The Approved Tool List row (implements Section 3)

Cursor (harness). Tier / account type: Team or business account, Privacy Mode enforced by admin | **Training disabled:** Yes, with Privacy Mode on (Exhibit 07) | **Memory off / scoped:** Vendor memory: none in this configuration. Working memory is files in the matter folder, local and inspectable | **Retention terms:** ZDR with all model providers; abuse-flag exception; covered models admin-opt-in with a 30-day safety window (Exhibits 07, 12) | **Permitted input tier:** Tier 2 | **Agentic capabilities:** Yes: reads, edits, runs tasks over the scoped folder. No autonomous filing, sending, or external transmission (policy Sections 4 and 5) | **Approved date:** [date of adoption] | **Notes:** Covered-model opt-in decision recorded here: [accepted / declined, date]. Codebase indexing decision recorded here: [on / off; see Section 3, step 5]

3. Required setup (implements Sections 3 and 4)

Steps 1 through 6 are the Administrator's; steps 7 through 9 are every user's.

1. **Buy team or business seats, not individual ones.** Privacy Mode can be enforced by a team or enterprise admin, and new team members inherit the team's setting (Exhibit 14). That inheritance is the enforcement mechanism; individual accounts leave the setting to each user's discretion.
2. **Enable Privacy Mode at the team level and verify per seat.** Record the date. Each user confirms the Privacy Mode indicator in their own settings at onboarding; the sixty-second check is part of training.
3. **Read the terms at signup and file them.** Pull the business-tier terms of service at signup and confirm they carry the data-use page's commitments; request the SOC 2 Type II report and the subprocessor list through trust.cursor.com (Exhibit 14) and file both with this document.
4. **Make the covered-model decision, deliberately.** Cursor blocks covered models until an administrator affirmatively accepts the 30-day safety-window terms (Exhibit 07 designates non-ZDR models; Exhibit 12 states the window). Decide once, in writing, on the row: either accept (knowing the window covers only traffic routed to those models) or decline and use non-covered models for client work. Do not let a deadline make this decision.
5. **Make the indexing decision.** If codebase/folder indexing is on, Cursor uploads content in chunks to compute embeddings; plaintext ceases to exist after the request, but embeddings and metadata (hashes, file names) may be stored (Exhibit 07). File names alone can identify a matter. Default for client work: indexing off for sensitive matters; where indexing is on, the folder's file naming follows the Placeholder Card discipline.
6. **Confirm non-ZDR models stay off.** Non-ZDR models are designated as such or require admin opt-in (Exhibit 07). The approved configuration does not opt in for client work.
7. **Users: one matter, one folder, one session.** The agent's reach is the folder, so the folder is scoped to the matter (policy Section 4, agentic access limits). No firm-wide folders, no "everything" workspaces. Substitute identifiers in the folder before work where the matter calls for it; substitution before transit is the harness's unique option.
8. **Users: nothing leaves without a human.** The agent does not file, send, sign, or communicate with anyone outside the firm; every outbound artifact gets attorney review (policy Section 5, agentic prohibition, verbatim duty).
9. **Users: complete training** on this document before first client use, including reading a session log so the record habit in Section 6 is concrete.

4. What may go in (implements Section 4)

- **Tier 1:** permitted.
- **Tier 2 (client documents and matter facts):** permitted in the approved configuration; this is the way of working where the policy's enterprise-configuration requirement is satisfied by architecture plus the ZDR terms, and where identifier substitution can happen before anything transits at all.
- **Tier 3:** prohibited absent informed client consent and Administrator approval. Note that the harness is the strongest candidate venue if consent ever makes Tier 3 work permissible, precisely because of the local footprint; that call remains the Administrator's, case by case.

This product's specific hazard is agency. The agent can read everything in the folder, including what you forgot was there; it can act quickly and plausibly, including wrongly. The controls are the two user rules above (scope the folder; nothing leaves without a human) plus the review discipline in Section 5. The Placeholder Card still applies to prompts; the folder-level substitution option supplements it, not replaces it.

5. Verification (implements Section 5)

Unchanged: the ladder, the statute check, the no-delegation rule. The harness helps rather than hinders here: it works over sources you put in the folder (closed-corpus grounding is the native workflow, not a workaround), and every edit is a visible diff. The rule stands anyway: the reviewing attorney examines the changes, not just the final text, before accepting them, and court-bound citations get the full four steps regardless of how grounded the drafting session felt.

6. The record (implements Sections 8 and 10)

The harness's structural advantage: when a session ends, the record already exists, locally. Session logs, instructions, and file versions sit in the matter folder, which is Section 10's documentation, the billing reconstruction record, and the troubleshooting record, all in one place, at zero vendor retention. The habit is preservation rather than export: the logs stay in the matter folder, are treated as attorney work product (policy Section 4), and are not pruned for tidiness.

Billing: seat subscriptions are overhead. If the firm adopts task-unit pricing (policy Section 8), the session logs are what make the per-task entries auditable.

7. Vendor facts and evidence grades

Privacy Mode: no training on Customer Data; ZDR agreements with all model providers; providers do not store or train. Source: Exhibit 07 | Grade: Published | Date: Retrieved 2026-08-07

Abuse-detector exception: flagged prompts may be stored for investigation. Source: Exhibit 07 | Grade: Published | Date: Same

Non-ZDR models designated or admin-opt-in. Source: Exhibit 07 | Grade: Published | Date: Same

Admin-enforced Privacy Mode, inherited by new team members. Source: Exhibit 14 | Grade: Published | Date: Same

SOC 2 Type II attestation available on request; annual penetration testing; subprocessor list on trust portal. Source: Exhibit 14 | Grade: Published; becomes document-verified when the firm pulls the report at signup | Date: Same

Indexing stores embeddings and file-name metadata; plaintext not retained; cached files encrypted with client-generated keys. Source: Exhibit 07 | Grade: Published | Date: Same

Covered-model 30-day safety window, admin-accepted, model-scoped. Source: Exhibits 07 and 12 | Grade: Published | Date: Same

Requests transit Cursor's backend even with a user-supplied API key. Source: Exhibit 07 | Grade: Published | Date: Same

Not established: the text of Cursor's ZDR agreements with providers (vendor-to-vendor contracts; inherently unpublished). **The ask for the vendor:** an enterprise or business agreement that restates the data-use page commitments as contract terms for the firm's account, plus the SOC 2 report and current subprocessor list (both available through the trust portal per Exhibit 14).

8. Re-vetting triggers (implements Section 12)

1. **Privacy Mode terms or scope change** on the data-use page (Exhibit 07's source).
2. **The provider list changes.** The page currently names SpaceXAI, OpenAI, and Anthropic; a new provider means new retention policies behind the ZDR umbrella.
3. **Covered-model designations expand** or the safety-window terms change (Exhibit 12's source page).

4. **Non-ZDR model handling changes** (designation or opt-in mechanics).
5. **Indexing or caching behavior changes.**
6. **Pricing model changes** (per-seat to metered; the policy's Section 3 vetting item 6 exhaustion questions get re-asked).
7. **Acquisition or material corporate change** at Anysphere.