Práxis

Chain of custody, intellectual property and data residency

Version 1.0. Last updated 6 September 2026.

This text is a starting draft and should be reviewed by legal counsel before launch.

This pack is the long form companion to our methodology page. It is written for legal, security and procurement reviewers who need the operational detail: how a single trajectory is tracked from creation to delivery, how intellectual property moves through the contractual chain, and which data residency configurations we can support. Every section here describes practice and process only. It states no benchmark results and names no customers.

1. Scope of this document

Práxis produces expert reasoning data: structured trajectories in which a verified professional works a synthetic case and records the decisions, the alternatives rejected and the checks performed. This document covers how that production is governed. It does not cover commercial terms, which live in the master services agreement, or the technical schema, which ships with each delivery report.

Where this document and a signed agreement differ, the signed agreement prevails. We update this pack whenever a control, a subprocessor or a residency option changes, and each version carries the version number and date shown above.

- Applies to all production work in the Práxis contributor network
- Applies to task design, review, audit and delivery
- Does not replace the master services agreement or a data processing agreement

2. Chain of custody in full

Custody is tracked as a sequence of recorded events attached to a single item identifier. The identifier is issued when a task instance is created and never reassigned. Every event that follows references that identifier, the actor who performed it, the software version in use and a server side timestamp. Nothing in the pipeline is edited in place: a correction is a new event that supersedes an earlier one, and the earlier one remains readable.

The stages below are mandatory. An item that has not passed each stage cannot enter a delivery batch, and the delivery tooling refuses batches containing incomplete custody records.

- Task authored: case identifier, author, domain, rubric version and difficulty calibration recorded

- Assignment: contributor identifier, registration status at assignment time and confidentiality confirmation recorded
- Capture: instrument session opened, autosave events written, elapsed working time recorded
- Submission: contributor submission sealed, no further contributor edits accepted
- Review: reviewer identifier, rubric version, criterion scores and written annotations recorded
- Adjudication: where reviewers disagree beyond tolerance, a senior reviewer decision is recorded with reasoning
- Audit sampling: sampled items marked with the auditor, the sampling rule and the outcome
- Batch assembly: item list, batch identifier, checksum and delivery report generated
- Delivery: recipient, transfer method, checksum acknowledgement and delivery timestamp recorded
- Retention or destruction: retention clock started, destruction event recorded when it expires

3. Custody records available to a customer

On request, and for any item in a delivered batch, we can produce a custody report showing the full event sequence for that item without exposing contributor identity. Contributors appear as stable pseudonymous identifiers so a reviewer can confirm that two items came from different people, or the same person, without receiving personal data.

If a regulator or an internal audit function requires identity disclosure, we can provide it under a separate written request, limited to the specific items in question and subject to the applicable data protection basis.

- Per item custody report with pseudonymous actor identifiers
- Batch checksum and delivery report for integrity verification
- Rubric version and criterion level review scores for each item
- Audit sampling outcome where the item was sampled

4. Intellectual property chain

The chain has exactly two links. Contributors assign all intellectual property in their work product to Práxis at the moment of submission, under a written agreement signed before their first assignment. Práxis assigns that intellectual property to the customer on delivery, under the commercial agreement. There is no third link, no retained contributor licence over delivered work and no collecting society in the path.

Task designs authored by Práxis for a specific customer engagement are treated as customer work product where the agreement says so. Task designs authored for our own general library remain ours, and we do not deliver library designs as exclusive material unless the agreement

states exclusivity explicitly.

Moral rights, where they cannot be assigned under applicable law, are waived to the maximum extent the law permits, and contributors agree not to assert them against the customer or against downstream users of a trained model.

- Contributor to Práxis: full assignment on submission, signed before first assignment
- Práxis to customer: assignment on delivery under the commercial agreement
- Warranty of original authorship from every contributor for every item
- Contractual prohibition on reproducing third party copyrighted text
- Explicit treatment of exclusive versus library task designs
- Waiver of unassignable moral rights to the extent the law allows

5. Data origin and third party material

Every production case is synthetic and authored by our task design team. Contributors are contractually prohibited from importing client files, patient records, case files, internal memoranda or any other third party confidential material, and they confirm that prohibition on each assignment before the instrument opens.

Review includes an origin check. A reviewer who recognises a case element as derived from a real identifiable matter, or as reproduced from a published work beyond permitted citation, rejects the item and the case is withdrawn from the library pending redesign. Rejections of this kind are logged as custody events like any other decision.

6. Data residency options

Residency is configured per engagement, before production starts. The default configuration keeps production data in a managed database region we select for latency and availability. Where a customer has jurisdictional requirements, we can pin production and delivery storage to a specific region, and we can restrict the human roles that may access data to people located in agreed jurisdictions.

Residency and access location are distinct controls. Pinning storage to a region does not by itself restrict who may open the data, so customers with strict requirements should specify both. We confirm the chosen configuration in writing and record it in the engagement file.

- Option A, Brazil residency: production, review and delivery storage in a Brazil region, contributor network and review staff in Brazil
- Option B, European Union residency: production and delivery storage in an European Union region, with review access restricted to agreed jurisdictions
- Option C, United States residency: production and delivery storage in a United States region for customers contracting with our Delaware entity
- Option D, customer controlled destination: delivery into storage the customer owns, in the customer's own region, with Práxis retaining only custody metadata

- Access location control: review and audit roles limited to named jurisdictions, independent of where storage sits
- Zero retention delivery: production copies destroyed on acceptance, with only custody metadata and checksums retained for audit

7. International transfer

Where a configuration requires personal data to move between jurisdictions, transfers rely on the mechanisms available under the applicable law: standard contractual clauses for transfers out of the European Economic Area, and the international transfer provisions of the LGPD for transfers out of Brazil. Production task content contains no third party personal data by construction, so transfer exposure is limited to contributor administration data.

Subprocessors are bound by written data processing terms that flow down the same obligations we accept. The current list is Supabase for managed database infrastructure and Resend for transactional email. We review the list whenever our stack changes and notify customers of additions where the agreement requires it.

8. Access control and confidentiality

Access is role based and scoped to a single engagement. Contributors see only the tasks assigned to them. Reviewers see only the batches they are assigned. Administrative access to production data is limited to named staff, granted for a stated purpose and revoked when the purpose ends. Every administrative session is attributable to an individual account.

Customer identity is not disclosed to contributors unless the customer agrees in writing. Task specifications are treated as customer confidential information, and we do not reuse a customer's specifications, rubrics or trajectories for another customer.

- Role based access with engagement level scoping
- Individual accounts only, no shared credentials
- Encryption in transit and at rest across the pipeline
- Least privilege administrative access, granted for a stated purpose
- Contributor confidentiality agreement signed before first assignment

9. Retention and destruction

Retention is set per engagement. The default is to retain production copies while the engagement is active and for a short acceptance window after final delivery, then destroy them and keep only custody metadata, checksums and review records for audit. Contributor administration data is kept for as long as the professional relationship lasts and afterwards only where a legal obligation requires it.

Destruction is an event in the custody log, so a customer can be shown when a copy was destroyed and by which process. Where a customer requires a certificate of destruction, we issue one referencing the batch identifier and checksum.

10. Delivery, provenance and acceptance

Batches ship as structured data in the agreed format, accompanied by a delivery report and a per item provenance record. The provenance record states the task version, the review trail, the audit status and the batch checksum. Batches are immutable once delivered: a correction ships as a new versioned batch that references the superseded items.

Acceptance is a defined step. The customer verifies the checksum, samples the batch against the agreed rubric and either accepts or returns findings within the window set in the agreement. Returned findings are worked as a remediation batch under the same custody rules.

- Structured delivery in the agreed schema and format
- Delivery report with batch composition and checksum
- Per item provenance record with review and audit trail
- Immutable batches, with corrections shipped as new versions
- Defined acceptance window and remediation path

11. Incident response

A suspected loss of confidentiality, integrity or availability affecting customer data is escalated internally on discovery, contained, and assessed for notification duties under the LGPD and the GDPR. Where the agreement sets a notification window we meet it, and where the law sets a shorter one the law governs.

Every incident produces a written record covering what happened, what data was in scope, what we did, and what changed afterwards. Customers affected receive that record.

12. Contact and requests under this pack

Requests for a custody report, a certificate of destruction, the current subprocessor list or a signed data processing agreement go to hello@praxis.ai. We answer procurement and security questionnaires directly, and we can execute a mutual confidentiality agreement before sharing internal policy documents.

Práxis is a Forward AI company. São Paulo, Brazil and Delaware, United States.

hello@praxis.ai