

SEAL Legal Runtime – Executive Overview

For General Counsel & Managing Partners

Public · v1.0 · December 2025

Table of Contents

1. WHAT SEAL IS (AND IS NOT).....	3
2. THE PROBLEM SEAL SOLVES.....	4
3. HOW SEAL WORKS IN LEGAL TERMS.....	5
3.1 Governance Anchors – The Questions SEAL Enforces.....	5
3.2 Sealed Outcomes – What You Actually See.....	5
4. SECURITY & ENFORCEMENT GUARANTEES (WHAT YOU CAN RELY ON).....	7
5. WHERE SEAL SITS IN YOUR STACK.....	9
6. LICENSING & ACCESS PATHWAYS.....	11
6.1 Three Access Modes.....	11
6.2 What You Get vs. What You Never Get.....	11
7. WHAT CHANGES FOR LAWYERS & STAFF.....	12
Day-to-day impact:.....	12
8. FREQUENTLY ASKED QUESTIONS.....	13
NEXT STEPS.....	14

1. WHAT SEAL IS (AND IS NOT)

SEAL Legal Runtime is a sealed enforcement layer that sits between your existing systems and any court, regulator, or counterparty.

- It **does not** replace your:
 - Governance / Risk / Compliance (GRC) program
 - Identity & access management (SSO/IdP)
 - Matter / case management or DMS
- It **does**:
 - Enforce your own rules at the point of filing or high-risk legal action
 - Require key governance conditions (role, vertical, motion, consent, authority) to be satisfied
 - Produce **sealed, audit-ready artifacts** for every approval, refusal, or override

Think of SEAL as a **governance gate** in front of your “file / submit / act” buttons for the workflows that call it. It never replaces your systems; it gates them and produces sealed evidence of what it did. When you wire it as the only path to a given action, every request on that path is evaluated under the same checks.

It simply answers:

“Given our policies and authorities, is this action allowed, refused, or requires supervision — and can we prove that later?”

2. THE PROBLEM SEAL SOLVES

Today, most firms and legal departments rely on:

- Case management, checklists, emails, and human review
- GRC and identity systems that sit **around** the workflow, not in front of it
- After-the-fact investigations when something slips through

That leaves recurring gaps:

- **Role & authority:** Was this the right lawyer (or agent) acting in scope?
- **Vertical & venue:** Was this filing even permitted in this court or practice area?
- **Consent & risk:** Did the client consent, and was the motion allowed under firm policy?
- **Proof:** If a bar complaint, malpractice claim, or regulatory inquiry arrives, can you show what your controls actually did at the moment of action?

SEAL closes that gap by moving governance from **policy documents and training** into a **sealed runtime checkpoint** that fires on every governed action, with its own evidentiary record.

3. HOW SEAL WORKS IN LEGAL TERMS

3.1 Governance Anchors – The Questions SEAL Enforces

Every decision SEAL makes is anchored to five declarations you already govern:

1. **Who is acting?** – Legal role (partner, attorney, paralegal, intern, AI agent, system account)
2. **Where are they acting?** – Legal environment / vertical (e.g., criminal defense, civil litigation, bankruptcy)
3. **What are they trying to do?** – Specific task or filing (e.g., motion for bail, motion to dismiss)
4. **How quickly?** – Turnaround / urgency (standard, expedited, emergency)
5. **Has the client consented?** – Consent gate as defined in your own policies

SEAL's approval is a **necessary pre-condition**, not a sufficient condition, for action: attorneys remain responsible for determining whether any filing or motion is strategically and ethically appropriate. It only answers:

“Is this combination of role + environment + motion + urgency + consent allowed under this tenant's rules and license?”

If those anchors are missing, inconsistent, or out of policy, the request is **refused with a sealed decision**, not “best-effort processed.”

Jurisdiction-specific content. For each practice area or jurisdiction, the firm (or its chosen advisors) is responsible for defining which roles, motions, venues, and consent conditions are permitted. Thinking OS™ does not supply or maintain jurisdiction-specific legal rules; SEAL enforces whatever the firm has adopted as its own policy.

Who owns the rules. Vertical policies and motion rules enforced by SEAL are authored and owned by the firm's legal leadership (or their designated advisors). Thinking OS™ does not determine which filings are lawful, advisable, or permitted in any jurisdiction. SEAL only enforces the firm's written governance rules and license scope, under the firm's supervision.

3.2 Sealed Outcomes – What You Actually See

For every governed request, you get:

- **Approval artifact** – “This was allowed”
- **Refusal artifact** – “This was blocked and here’s why”
- **Override artifact** – “This would have been refused, but an authorized supervisor overrode it with proof”

Each artifact includes:

- A unique **decision / trace ID**
- A **cryptographic hash** (tamper-evident)
- The anchors in force (role, vertical, motion, turnaround, consent)
- A **code family** (e.g., consent, identity, license, vertical, motion policy, safety) and human-readable rationale

These artifacts are designed to be:

- Attached to matters
- Produced in discovery or bar / regulator packets
- Queried by risk & compliance (“show all consent-gate refusals this quarter”)

Internal enforcement logic, models, and prompts are never exposed. Regulators and auditors test SEAL by sending scenarios and reading outputs, not by inspecting code.

Interpreter & classification safeguards. Where SEAL uses interpreter or classification logic to normalize inputs or categorize risk, those components are configured to fail closed: uncertainty results in a structured refusal or advisory, not a silent approval. Any external AI services used in these steps operate under “no training on customer data” and are treated as subprocessors in client DPAs.

4. SECURITY & ENFORCEMENT GUARANTEES (WHAT YOU CAN RELY ON)

From a GC/MP perspective, SEAL makes a small set of hard promises:

1. Governance gate for wired workflows

- For any workflow wired to SEAL as the decision point, all approved filings pass through the same pipeline: intake → checks → decision → record.
- Inside the runtime, there is no alternate path that returns an approval without running identity, consent, and policy checks.
- Whether a given filing must pass through SEAL is controlled by how your own systems are integrated (e.g., wiring SEAL as the only route to “file / submit / act” for that workflow).

2. Fail-closed by design

- Ambiguous identity, unknown roles, unlicensed verticals, missing consent, broken evidence/authority, malformed payloads, or unreachable providers all result in **sealed refusals**, not silent passes or 5xx errors.

3. Client-owned identity & GRC

- Roles and permissions come from your IdP / SSO and GRC; SEAL never invents or stores a “shadow org chart.”
- Legal roles are mapped from your groups; unknown or unmapped roles are refused.

4. Immutable, auditable records

- Artifacts are written to **client-owned, append-only audit storage** with write-once posture if you choose.
- SEAL never edits artifacts in place; corrections create new decisions linked by trace ID.

5. Tenant isolation

- Each `client_id` has its own governance perimeter; no cross-tenant data mixing, analytics, or shared dashboards.

6. Minimal surface area

- Only two things are exposed: the intake API and sealed artifacts / events. No “rule inspector,” no debug console, no model knobs.


7. Data protection & access control

- Thinking OS personnel **cannot access client artifacts for product experimentation or model tuning**; access is limited to support and incident response under strict logging.
- SEAL artifacts are **encrypted in transit** and at rest in standard cloud infrastructure.
- Artifacts are **never used to train third-party models** or shared for analytics.
- Access to client artifacts by Thinking OS™ personnel is restricted to a small, audited support and incident-response group under role-based access control.
- Firms may **direct artifacts into their own storage perimeter** and set their own retention policies.

For you, this means the worst-case scenario is a **documented, explainable refusal**, not an invisible bypass.

When we detect a defect or misconfiguration in the runtime that could materially affect approvals or refusals, we notify affected clients and work with their risk and IT teams on a remediation plan. Firms, in turn, remain responsible for monitoring their own queues and treating unexpected refusal patterns as signals to investigate rather than assuming “the system is always right.”

5. WHERE SEAL SITS IN YOUR STACK



SEAL sits as a sealed enforcement layer between the firm's own systems (GRC, identity, matter) and any downstream action or automation. It never replaces those systems; it gates them and produces sealed audit evidence.

- **Inputs:**
 - Your GRC & policies
 - Your IdM / SSO / org chart
 - (Optionally) your matter / case systems and terminology registry
- **SEAL Runtime:**
 - Verifies identity, license, vertical, motion, consent, evidence, authority
 - Produces a sealed decision + artifacts
- **Outputs:**
 - Sealed artifacts into your audit storage
 - A signed decision event into **your** router / notification service

- From there, your systems decide who sees what (attorneys, docketing, GC, audit, etc.)

SEAL never becomes your matter system or your routing engine; it simply inserts a sealed governance checkpoint into the workflows you already own.

SEAL Legal Runtime – Exec Overview v1.0 (2025-12)

6. LICENSING & ACCESS PATHWAYS

6.1 Three Access Modes

1. **Governed Runtime Exposure (pre-license)**
 - For bar committees, ethics boards, GC, and risk teams.
 - Vendor baselines only; **no connection** to your IdM/GRC.
 - Tight scope and duration; used to “watch the gate work” with your scenarios under seal.
2. **Sealed / Integrated Pilot (pre-license, client-wired)**
 - SEAL is wired to your identity, groups, and (optionally) GRC and matter systems.
 - Your policies drive decisions; vendor baselines remain as a safety net while configuration is hardened.
 - Used to confirm that approvals/refusals line up with your real rules and ethics expectations.
3. **Command-Layer License (production)**
 - Per-tenant license; covers **all shipped legal verticals** and **unlimited reviewers / filings**.
 - Identity and authority come solely from your IdM and GRC; SEAL enforces strictly, fail-closed.
 - Vendor baselines are only used where explicitly agreed as safety rails.

6.2 What You Get vs. What You Never Get

You get:	You never get (by design):
• A dedicated tenant profile (identity posture, routing, legal scope, license block) • Right to route production filings through SEAL • Production use of sealed artifacts and decision events for your matters	• Access to internal algorithms, prompts, or model configurations • Capacity to repurpose SEAL as a general-purpose AI governance layer • The right to host SEAL internals inside your own environment

The license gives you **access to the judgment perimeter — not the machinery inside it.**

7. WHAT CHANGES FOR LAWYERS & STAFF

Day-to-day impact:

- Lawyers and staff work in **the same systems** (DMS, case management, intake tools)
- Before a high-risk filing or action goes out, those systems call SEAL with the anchors
- The team receives one of three things:
 - **Approved** – proceed; artifact attached to the matter
 - **Refused (with code & rationale)** – fix the issue or escalate as per firm policy
 - **Advisory / requires override** – a partner or authorized supervisor decides, with proof

Even when SEAL returns an approval, **lawyers and supervisors must still exercise independent judgment**. SEAL confirms that the request meets the firm's configured criteria; it **does not certify outcome quality, merit, or ethics**.

There is no "SEAL dashboard" for attorneys to learn. SEAL shows up as:

- A consistent approval / refusal message
- A link to a sealed artifact when needed
- A predictable set of codes and reasons that risk/GC can query

Example:

An associate clicks 'File motion for bail' in your case system. Before filing, the system calls SEAL with who/what/where/consent. SEAL either returns a sealed approval (artifact attached to the matter) or a refusal with a code like SEAL-CONSENT-XXX and a short explanation. The associate fixes it or escalates per your policy.

8. FREQUENTLY ASKED QUESTIONS

Q1. Can this help us with bar, regulator, or malpractice scrutiny?

Yes. Every governed decision has a sealed artifact with anchors, rationale, hashes, and trace IDs. That gives you contemporaneous evidence of **what your governance layer did** at the moment of action, even years later.

Q2. Does SEAL replace our GRC or IdM?

No. You must retain — and continue to improve — your own policies, role definitions, and identity systems. SEAL's value is that it **enforces those decisions at runtime** and proves it, without ever becoming your system of record.

Q3. What if SEAL makes a mistake?

SEAL never invents policy; it enforces the configuration and policies you supply. If an outcome looks wrong, the artifact shows exactly which anchors and policy posture were in force so you can correct your own configuration or rules. Attorneys and supervisors are expected to treat SEAL's decisions as governance signals, not as a substitute for legal judgment; approvals are pre-conditions, not guarantees that a filing is strategically or ethically sound.

Q4. What happens if SEAL or the network is down?

The design philosophy is: **no silent degradation**. If SEAL cannot safely evaluate a request, it returns a sealed refusal / error family rather than allowing ungoverned filings to continue unseen. Operational SLAs and fallbacks can be addressed in the commercial agreement and your own business-continuity plans. Firms are expected to maintain human checks for critical deadlines; SEAL refusals are treated as incident signals, not automatic safe-harbor.

Q5. Can regulators see how it works inside?

No — and that is intentional. Regulators and auditors test SEAL the same way you do: via standard payloads and sealed outputs. The artifacts and refusal/approval codes are designed so they can map them to ABA themes, ISO/FRCP obligations, and local rules without needing or getting access to internal logic.

Q6. Is any of our data used to train models?

No. We explicitly **reject using client artifacts for training** or **exposing prompts/model internals**, and we contractually require the same from any subprocessors we use. Thinking OS's IP and your IP stay sealed; the only shared surface is structured, immutable outputs.

Q7. How do we escalate if we think SEAL is misbehaving?

Firms treat unusual refusal patterns or suspected misconfigurations as operational incidents: pause reliance on affected flows where necessary, review sealed artifacts, and escalate to Thinking OS™ support with trace IDs. SEAL's design makes systemic issues visible as clusters of structured refusals, rather than hidden silent passes.

Q8. How to describe SEAL externally?

"Our firm uses a sealed governance runtime (Thinking OS™ SEAL) to enforce internal approval rules before certain filings or actions. This runtime does not provide legal advice or draft documents. It helps us apply our own policies consistently; attorneys remain fully responsible for all decisions and filings." This language emphasizes that SEAL is a governance control under attorney supervision, not a substitute for legal judgment.

NEXT STEPS

If you'd like to see SEAL applied to your own scenarios, the next step is a governed runtime exposure session under NDA, where your risk team can send real patterns and review sealed outputs.