

BenchSlap and the Option.Black Portfolio

A platform overview: nine verification-first tools, four hostname-routed surfaces, eight invented subsystems, one founding architectural commitment.

Author: Richard L. Sanders · Utah Bar #15728 · Founder, Option.Black · Author, *The Phoenix Conspiracy* series **Role:** Sole inventor, architect, and implementer of every subsystem named in this paper **Version:** 1.0 — 2026-05-13 **Source:** the public repository at <<https://github.com/Benchslap/Benchslap>>; production deployment at benchslap.pro, benchslap.com, benchslap.app, 4all.help, option.black, optio.black **Companion papers:** [01-deterministic-closed-corpus-verification.md](#) (AEGIS), [02-opinion-harvester.md](#), [03-commission-inquest.md](#), [04-chomper-evidence-ingestion.md](#), [05-v4-advocacy-algorithm.md](#), [06-citation-gravity.md](#), [07-vox-omnium.md](#), [08-aegis-prime-structural-verification.md](#), [09-benchslap-tool-catalog.md](#)

Abstract

The commercial AI-assisted legal-research market consists of products that wrap a language model in a chat interface and call it a tool. The wrapper sits between the user and the model, but the wrapper cannot tell whether the model fabricated the citation, misstated the holding, applied the wrong burden of proof, or chained a non-sequitur from real law to real fact. The user is told to verify. The user does not.

I built a different category of product. The BenchSlap platform is a **verification-first system** in which every legal claim, every citation, every fact, and every logical inference passes through deterministic gates before it reaches the user. The model is allowed to produce drafts; deterministic infrastructure decides whether those drafts may ship.

The platform consists of:

- **Nine AI tools** (Reckoner, Drafter, Auditor, Preprocessor, Strategist, Advisor, Commission Inquest, Option.Black Council, and the V4 Engine that gates them all)
- **Four hostname-routed surfaces** (benchslap.pro for attorneys, benchslap.com for pro-se, benchslap.app for Labs/sandbox, 4all.help for the corrections population)
- **Eight invented subsystems** (AEGIS content-hash pinning, AEGIS PRIME structural verification, the V4 Advocacy Algorithm, the Cascading Opinion Harvester, Citation

Gravity, Chomper evidence ingestion, Commission Inquest, Vox Omnium translation)

- **One founding architectural commitment** — that information must not be destroyed, that retrieval-augmented generation is unacceptable for high-stakes legal work, and that verification must be deterministic-first.

This paper is the umbrella. Each subsystem named in the eight-item list has its own white paper or USPTO provisional, identified by hyperlink in the relevant section.

1 • The founding commitment

As soon as you destroy information you have to accept it is gone or, as the model is trained to do, reproduce it by a bunch of guessing that may be statistically informed but does not replace actual truth.

That statement is the architectural thesis. Every subsystem on this platform proceeds from it.

The argument is concrete and load-bearing. Commercial RAG (retrieval-augmented generation) destroys information through (a) **chunking** that splits the source by character count rather than by semantic boundary, (b) **embedding compression** that projects natural language into a fixed-dimension vector space with lossy reconstruction, and (c) **top-K retrieval** that returns the K nearest neighbors and discards the (K+1)th context the synthesis layer needed. The synthesis layer then reconstructs the missing information by inference — which is exactly the failure mode the user came to a non-hallucinating system to escape.

The platform's architectural alternative is **closed-corpus verification with content-hash pinning**. The corpus is fully indexed. Every opinion has a content hash. Every citation is verified by substring or shingle containment against the corpus's stored text — not by similarity in an embedding space, but by literal text match against the corpus text the model claims to be citing. Tamper is detected. Hallucination is foreclosed at the gate, not at the disclaimer.

2 • The four hostname-routed surfaces

A single Node.js application serves four distinct surfaces, distinguished by `req.host`:

Hostname	Surface	Audience	Model tier
benchslap.pro	BenchSlap PRO	Practicing attorneys, paralegals	PRO — Claude Opus 4.7, Gemini 3 Pro, GPT-5.2 Pro
benchslap.com	BenchSlap Free	Pro-se litigants, general public	Free — Gemini 3 Flash (cost-optimized kernel), parallel-audit by Gemini Flash + GPT-5.4 Mini
benchslap.app	Bench Slap Labs	Internal sandbox for new features	PRO-tier (replica of PRO)
4all.help	4All.Help (corrections-focused front)	Incarcerated and underserved	PRO-tier with plain-language response register
option.black	Option.Black umbrella brand	The portfolio's public face	n/a (umbrella)
optio.black	Conclave apex	6-model deliberation forum	n/a (forum)

The audience-aware register shift on **4all.help** is enforced at the prompt-tier level: same backend tools, same corpus, but the system prompt directs the model to explain terms inline ("due process" — which means fair treatment under the law") and to give actionable instructions rather than attorney-grade structural advice. Same engine, different voice; no separate codebase.

The four-tier hostname routing keeps mission separation visible at the URL and forces the brand identity, color palette, and audience register to stay coherent end-to-end. There is no cross-domain user confusion about which tool they are looking at.

3 · The nine AI tools

Each tool is a route handler in **routes/tools.js** (or **routes/tools/option-black.js** for the Council) that funnels the user's request through the same verification stack before, during, and after model invocation.

3.1 · Reckoner (Consigliere)

Multi-turn chat assistant with silo context (case-specific working memory). Every cited authority verifies through the AEGIS-pinned corpus before the response stream completes. Tier: Free + PRO. Endpoint: **POST /api/chat**.

3.2 · Drafter

Generates new legal documents AND improves existing ones. Endpoints: `POST /api/draft` (generate), `POST /api/draft/improve` (improve existing), `POST /api/draft/chain` (multi-step pipeline). The legacy `POST /api/fix` is a backward-compatibility alias for `improve`. The same V4 verification stack runs over every Drafter output — a drafted brief that contains a fabricated citation fails the post-stream gate and never reaches the user.

3.3 · Auditor

Verifies an existing document and produces a strength report with GREEN/YELLOW/RED/BLACK color-coded flags per V4 truth-table cell. Three rounds of Auditor against the V3 Hymas motions (Hymas v. Hymas, 20251426-CA) produced 54-of-54 + 52-of-52 case citations verified, 0 repudiated, 0 AEGIS hash failures, with two reported "mismatch" warnings traced to Auditor name-extraction false positives (motion cites were correct — verified by direct corpus lookup against `2012 UT 79` = Utah DOT v. FPA West Point, `2022 UT App 47` = Knowles v. Knowles).

3.4 · Preprocessor

Seven-phase document-classification pipeline that builds an authority index (every cited case) and a fact freezer (every factual assertion frozen with a source pin) for any uploaded document. Feeds the downstream tools' working memory. Wired into the live upload path; the related Chomper evidence-ingestion engine (see [04-chomper-evidence-ingestion.md](#)) is the universal-modality extension currently developed as standalone infrastructure.

3.5 · Strategist (PRO only)

Quick case-strength analysis. Score (0–100), top three arguments, risk overview. Sub-2-minute response. Endpoint: `POST /api/strategy`. Built for the "before-I-take-this-case" sanity check.

3.6 · Advisor v2.0 (PRO only)

Deep four-layer cognitive processing (Parse → Extract → Analyze → Advise) with V4 Advocacy Analytics, Totem memory (Hot/Warm/Cold ring-based fact retention), and opponent-pattern detection across 30 known patterns in 8 categories (Volume, Concealment, Delay, Abuse, Intimidation, Deception, Discovery, Strategic). Endpoint: `POST /api/advisor`. Built for theory-of-the-case development and trial preparation.

3.7 · Commission Inquest (PRO only)

Proprietary multi-AI deliberation for ad-hoc design / code / strategy decisions. Convenes flagship models (Claude Opus 4.7, Gemini 3 Pro, DeepSeek V3, with weighted voting at 0.40 / 0.35 / 0.25) and resolves to UNANIMOUS / MAJORITY / SPLIT consensus modes. Endpoint: `POST /api/inquest` . Full architecture in `03-commission-inquest.md` .

3.8 · Option.Black Council (top-tier only)

Six-model conclave (the inquest's bigger sibling) hosted at `optio.black` with a benchslap-session-token bridge. The Council adds three additional models to the inquest's three for the highest-stakes deliberations (capital-case strategy, novel constitutional questions, drafting law-changing motions). Endpoint: `POST /api/tools/option-black` .

3.9 · The V4 Engine

Not a tool the user picks; it is the truth-table verification layer every other tool routes through. Three-bit deterministic classification over (Law × Fact × Logic). Nine cells, four color flags, eight parallel calculi. Full architecture in `05-v4-advocacy-algorithm.md` .

4 · The eight invented subsystems

The nine tools above are the user-facing surface. Underneath them is a verification stack of eight invented subsystems, each of which is independently novel and several of which have or will have USPTO provisional protection.

4.1 · AEGIS — Closed-Corpus Verification

Content-hash pinning of every authority. Every row in `opinion_library` , `rules_library` , and `court_rules` has a `content_hash` (SHA-256 over the normalized opinion text) and a `content_shingles` jsonb array (5-gram containment signature). 100% coverage on **3.2 million opinions + 1.9 million rules + 72 court rules**. Tamper is detected. Hallucinated holdings are blocked at the post-stream gate. Full paper: `01-deterministic-closed-corpus-verification.md` . USPTO provisional filed.

4.2 · AEGIS PRIME — Structural Verification

The closed-corpus extension. Pre-extracts structured facts (disposition: AFFIRMED / REVERSED / VACATED / etc., panel vote: majority / concurrence / dissent, holding-binding-weight: BINDING / PERSUASIVE / DICTA / REJECTED / SUPERSEDED, case-treatment graph: OVERRULED / SUPERSEDED / DISTINGUISHED) for every Utah case at ingest time. Hard-blocks `DISPOSITION_MISMATCH` (AI says "affirmed" but case was reversed), `ATtribution_MISMATCH` (AI quotes dissent as majority), `BINDING_WEIGHT_MISMATCH` (AI frames dicta as binding holding), and `SUPERSEDED_CASE` .

~8,000+ Utah dispositions extracted, 129 overruled cases, 238 treatment-graph edges. Documented in [docs/AEGIS_PRIME_WHITE_PAPER.md](#).

4.3 · The V4 Advocacy Algorithm

The deterministic-first truth-table verification engine. Three-bit classification, four-gate decomposition, eight parallel calculi (deterministic logic, Dempster-Shafer belief functions, Bayesian epistemological layer, defeasibility classifier, burden-of-proof evaluator, Toulmin warrant parser, argumentation framework with dialectical tracker, multi-model peer review). The deterministic floor is absolute — probabilistic and neural layers may refine but never override the deterministic verdict. Full paper: [05-v4-advocacy-algorithm.md](#).

4.4 · The Cascading Opinion Harvester

The .gov-direct, no-aggregator, weakest-first opinion-corpus pipeline. Harvests from every U.S. state court's public [.gov](#) website (all 50 states + DC + 13 federal circuits + 5 territories + NJ Appellate Division = 70 STATE_SOURCES). No commercial aggregator dependency. No CourtListener single-point-of-failure. As of paper date: **3.2 million opinions** with continuous ingest. Full paper: [02-opinion-harvester.md](#). USPTO provisional filed.

4.5 · Citation Gravity

The closed-corpus self-audit. Discovers the bedrock authority the corpus is missing by counting which citations the corpus's own opinions cite that the corpus does not itself contain. Memory-bounded SQL-aggregated implementation runs against a 3.2-million-row corpus inside a 4 GB heap. Closed-loop coupled with the harvester via [landmark-harvest.js](#), so the corpus tells the harvester what to harvest next. Full paper: [06-citation-gravity.md](#).

4.6 · Chomper — Universal Evidence Ingestion

One ingestion pipe for every common evidentiary file type (PDF, image, audio planned, video planned). Quantitative four-factor confidence formula $C = S \times (1 - A) \times R \times E$ on every observation. Extract-and-destroy provenance — originals zeroed in RAM before response returns. Full paper: [04-chomper-evidence-ingestion.md](#). USPTO provisional filed.

4.7 · Commission Inquest

The multi-AI deliberation engine that produces UNANIMOUS / MAJORITY / SPLIT consensus across flagship models for high-stakes decisions. Weighted voting at architecturally-tuned weights, three quorum modes (standard / strict / expedited), per-inquest audit log. Full paper: [03-commission-inquest.md](#). USPTO provisional filed.

4.8 · Vox Omnium Prime

The multi-model verification cascade for long-form literary and legal translation. Eight-pass internal pipeline within Pass 1 (extract / segment / map / translate / verify / refine / cohere / naturalize), three-pass cascade at the architecture level (translator / evaluator / tribunal). Author-supplied glossary, voice profile, and series bible as load-bearing first-class inputs. Cross-chunk entity seeding producing 99.4% canonical-rendering rate against DeepL's 31% on a 695,461-character literary corpus. Five-flagship-judge blind panel produced unanimous superior verdict. Full paper: [07-vox-omnium.md](#).

5 · Live corpus statistics


As of 2026-05-13:


- **3,206,590 opinions** in `opinion_library` (continuous ingest)
- **1,914,203 rules** in `rules_library` (continuous ingest)
- **72 court rules** in `court_rules`
- **100% AEGIS content-hash coverage** across all three tables
- **70 STATE_SOURCES entries** in the active harvester (50 states + DC + 13 federal circuits + 5 territories + NJ_APP)
- **133 distinct court codes** represented in the corpus
- **8 backup layers**: USB primary (exfatdrive), USB secondary (Untitled), DigitalOcean managed-PG daily snapshots, GitHub, plus continuous-cadence schedules (15-min repo rsync, 1-hr verification, 6-hr full pg_dump)

Live count at <https://benchslap.pro/api/corpus/stats>.

6 · The verification stack at a glance

For any user request reaching any of the nine tools, the request travels through:





There is no path through this stack that bypasses step 5, 6, or 7. The wrapper architecture is not "model with disclaimer" — it is "model with deterministic verifier downstream that refuses to ship what it cannot verify."

7 · The portfolio's identity

The eight invented subsystems plus the nine tools plus the four hostname-routed surfaces are not a collection of features. They are one platform that takes one architectural commitment and follows it through every layer:

Information is preserved. Verification is deterministic. Hallucination is foreclosed at the gate, not at the disclaimer.

This is the architectural commitment that distinguishes the platform from every commercial AI-assisted legal-research tool currently in market. The commercial tools cannot retrofit this commitment because their architectures are built on the RAG primitives that destroy the information the verifier would need to verify against. The retrofit is not an engineering task; it is an architectural rebuild. The portfolio's value is that the rebuild is already done.

8 · Document index

- [01 — Deterministic Closed-Corpus Verification \(AEGIS\)](#)
- [02 — The Cascading Opinion Harvester](#)
- [03 — Commission Inquest](#)
- [04 — Chomper Evidence Ingestion](#)
- [05 — The V4 Advocacy Authority Algorithm](#)
- [06 — Citation Gravity](#)
- [07 — Vox Omnium Prime](#)
- [08 — AEGIS PRIME — Structural Verification](#)
- [09 — The BenchSlap Tool Catalog](#)
- 00 — *This document* (platform overview)

Each subsystem paper stands alone and can be read without the others. Each names its inventor (the same one in every case), cites real file paths in the public repository, and is anchored by USPTO provisional protection or live code.

© Richard L. Sanders 2026. All rights reserved except those granted by the project's open license. BenchSlap™, Option.Black™, 4All.Help™, AEGIS™, AEGIS PRIME™, V4 Engine™, Citation Gravity™, Chomper™, Commission Inquest™, Cascading Opinion Harvester™, and Vox Omnium Prime™ are unregistered trademarks of the author.