

Public Defense AI Workflow Pilot

Records, Mitigation, and Attorney Review

Purpose

Public defense teams routinely face discovery, medical records, mitigation history, jail calls, transcripts, and client communications at a scale that exceeds available staff time. The useful question is not whether AI should replace legal judgment. It should not. The practical question is whether attorney-supervised AI workflows can reduce administrative burden while preserving confidentiality, source control, and human review.

This pilot is a review-ready proposal based on federal defense internship work, legal AI adoption research, and an empirical legal-data thesis. It is designed for offices that need small, auditable workflows before considering broader institutional adoption.

Pilot Scope

The pilot covers three attorney-supervised workflows:

- 1. Record triage and chronology**
 - Convert large record sets into source-linked timelines.
 - Flag medical, mitigation, and factual themes for attorney review.
 - Preserve citations to source pages or excerpts wherever possible.
- 2. Mitigation support**
 - Help staff organize life-history, treatment, family, trauma, education, employment, and reentry materials.
 - Draft structured issue maps, not final advocacy.
 - Keep attorney judgment central for narrative, legal theory, and client dignity.
- 3. Internal workflow documentation**
 - Produce office-facing checklists, risk notes, training examples, and review rubrics.
 - Document when AI is useful, when it is not, and which tasks require complete human handling.

Governance Principles

- **Attorney review is mandatory.** AI output is treated as a draft, index, or triage aid.
- **Confidentiality controls come first.** No live confidential data enters tools without office-approved privacy, retention, and access controls.
- **Source traceability is required.** Workflows should prefer outputs with page, file, excerpt, or record references.
- **No black-box conclusions.** The system should surface evidence, conflicts, and uncertainty rather than assert case strategy.
- **Bias and hallucination are operational risks.** Review checklists should include citation checks, source mismatch checks, omission review, and adversarial verification.

30-Day Deliverables

- Workflow map for one records-heavy defense task.
- Sanitized sample templates for chronology, record summary, and mitigation issue map.
- Attorney review checklist for citations, omissions, confidentiality, and judgment-sensitive outputs.
- Pilot evaluation rubric covering time saved, error modes, user trust, and review burden.
- Short implementation memo recommending whether to expand, pause, or redesign the workflow.

Example Review Packet Structure

The pilot output should be a review packet, not a free-form chatbot transcript. A sanitized template can use this structure:

Packet field	Review purpose
Source reference	File name or record label approved for the pilot, page range, date, and excerpt location.
Extracted fact	One narrow factual claim stated without legal conclusion.
Chronology row	Date, event type, people or institutions involved, and confidence level.

Packet field	Review purpose
Attorney check	Confirmed, needs source review, conflicting source, omitted context, or do not use.
Risk note	Confidentiality concern, hallucination risk, missing source, bias risk, or client-dignity concern.

This structure keeps the system focused on evidence location and review burden. It avoids strategy recommendations, legal conclusions, or unsupervised advocacy.

Evaluation Questions

The pilot should not be judged by whether AI output sounds polished. It should be judged by whether the workflow improves attorney review:

- Did the workflow reduce time spent locating relevant facts without increasing correction burden?
- Did the output preserve enough source references for fast verification?
- Did attorneys find the summaries useful for issue spotting, or did they create new review risk?
- Which tasks remained inappropriate for AI assistance even under supervision?
- Did staff understand the confidentiality boundary well enough to use the workflow safely?

Suggested Operating Model

A small pilot should start with a limited participant group: one supervising attorney, one attorney or mitigation specialist handling the records-heavy task, one technically fluent operator, and one reviewer responsible for checking citations, omissions, and output quality. The pilot should use de-identified or office-approved materials first, then expand only if the office is satisfied with privacy controls, review burden, and the quality of source-linked output.

The final product is not a chatbot. It is an attorney-reviewed workflow packet: inputs, allowed tools, prohibited uses, templates, review rubric, examples, failure modes, and a recommendation on whether the office should expand, revise, or stop the experiment.

Why This Matters

The technology gap in criminal defense is not just a matter of convenience. Resource asymmetry affects how quickly teams can understand records, communicate with clients, prepare mitigation, and respond to complex fact patterns. A careful pilot can improve institutional capacity without pretending that AI can replace lawyers, investigators, mitigation specialists, or client-centered judgment.

What This Demonstrates

This proposal demonstrates practical legal AI adoption judgment: risk-aware tool evaluation, workflow design, source-linked outputs, public-sector constraints, and the ability to translate technical systems into attorney-facing implementation.