



Suhas Bhairav

AI Legal Knowledge Assistant Blueprint

A practical, citation-first AI for internal legal knowledge with source controls and disclaimers.

Deliver a compliant, citation-backed internal legal knowledge assistant with clear boundaries and auditable controls.

Suhas Bhairav

suhasbhairav.com

Citation-based Knowledge Assistant

This page defines what to automate in legal workflows while preserving human oversight. It distinguishes automatable tasks, tasks needing review, and data/tool prerequisites.

Key move **Begin with a narrow use-case to prove citation fidelity.**

PRACTICAL CHECKLIST

- 01** Automate retrieval of authorities and case-law summaries with verifiable citations.
- 02** Automate drafting of memos and routine responses with source links.
- 03** Automate standard contract clause lookup and redline suggestions for common templates.
- 04** Keep human review for legal conclusions, risk judgments, and strategic advice.
- 05** Require source-of-truth verification and mandatory disclaimers before sharing outputs externally.
- 06** Store and version sources and outputs in a controlled repository.

Suhas Bhairav AI Resource

Deliver a compliant, citation-backed internal legal knowledge assistant with clear boundaries and auditable controls.

Tools, Data, Roles

This page maps data sources, tools, and roles required to operate a citation-based legal knowledge assistant. It covers data provenance, confidentiality boundaries, access controls, and versioning.

Key move Document ownership: assign a data steward per knowledge domain.

PRACTICAL CHECKLIST

- 01** Source-of-truth databases: statutes, regulations, case law, and official guidance.
- 02** Citation management: consistent formatting and linking to authors, opinions, and sections.
- 03** Confidentiality boundaries: data classification, access controls, and leakage protections.
- 04** Versioning: outputs and sources encoded with timestamps and provenance.
- 05** Permissions: role-based access, least privilege, and approval workflows.
- 06** Tooling: retrieval-augmented generation, vector store, and audit logs.
- 07** Human-in-the-loop: human review before finalizing outputs.

Suhas Bhairav AI Resource

Deliver a compliant, citation-backed internal legal knowledge assistant with clear boundaries and auditable controls.

Adoption Plan

Provides a pragmatic, week-by-week plan to pilot and scale the knowledge assistant. Emphasizes governance, approvals, and measurable outcomes.

Key move Assign a knowledge steward and governance lead from day one.

PRACTICAL CHECKLIST

- 01** Week 1: deploy a read-only FAQ assistant for internal policies with citation links.
- 02** Week 2: enable contract clause lookup and redline summaries for standard templates.
- 03** Week 3: integrate with document management, enable versioning, and restrict external sharing.
- 04** Week 4: establish human-review queue and escalation to counsel for high-risk outputs.
- 05** Week 5: implement access controls, role-based permissions, and audit logging.
- 06** Week 6: run a privacy and confidentiality audit; adjust data-handling policies.
- 07** Ongoing: track KPIs, governance approvals, and owner sign-offs for outputs.

Suhas Bhairav AI Resource

Deliver a compliant, citation-backed internal legal knowledge assistant with clear boundaries and auditable controls.

ABOUT THE CREATOR

Suhas Bhairav

I build AI systems that combine distributed architecture, retrieval, knowledge graphs, security, and practical workflow design. My focus is AI that can be used repeatedly by real teams: understandable, reliable, and grounded in the work people already do.

EXPERTISE

- Principal systems architecture and applied AI engineering
- Production-grade AI systems, agentic workflows, RAG, and knowledge graphs
- Full-stack workflow applications across AWS, Google Cloud, Azure, on-premise, and hybrid environments
- Practical AI implementation patterns for automotive, manufacturing, finance, logistics, cybersecurity, sales, and customer operations
- Security, governance, observability, approvals, and failure-mode thinking for operational AI systems

WEBSITE

suhasbhairav.com

EMAIL

suhasbhairav@gmail.com