

AI Solutions Overview

Hackathon for a Better World 2026

Harnessing AI Safely & Effectively

The following 14 solutions were developed in response to the Harnessing AI Safely & Effectively challenge statement of Hackathon for a Better World 2026. The summaries reflect the proposals as submitted and provide an overview of the issues identified and approaches proposed by participating teams. They do not constitute recommendations or endorsements by SAL.

Overview of Solutions

No.	Team	Solution
1	Lex et Machina (Honorary Team)	Counsel Companion
2	Misfits	Margin
3	Chubbles	CiteCheck SG
4	Team Crank	From Principles to Practice
5	Jam & Raph	Form Filler
6	WongP	Clawset
7	Red Pulse	Agent Red
8	WP-KIT	RADAR
9	TSMP Disputes	AI Twinning
10	NextGen Legal	LawLock
11	Condition Precedent	Docket: The First Matters Companion
12	USP Law	Offline Redaction Tool
13	OptiNUS Prime	Raising the Bar with Artificial Intelligence
14	ANGUS	Prompt-Building Tool

1

Lex et Machina (Honorary Team) Counsel Companion

Issue addressed: Clients increasingly have access to AI-enabled self-help tools, while lawyers must continue to differentiate themselves through judgement, trust and a deep understanding of each client's circumstances. Lawyers also need practical ways to maintain client context across matters and identify emerging risks without weakening the human relationship at the centre of legal advice.

Proposed approach: Counsel Companion is an AI-enabled client-centred support tool that maintains a continuously updated, client-specific context repository using authorised information such as matter intake, client correspondence, historic advice and relevant external developments. At intake, it identifies missing information and helps clarify the client's circumstances; during the relationship, it monitors relevant developments and surfaces tailored risk alerts; and when counsel prepares advice, it reviews the work product against the accumulated client context and suggests ways to improve relevance. The team's core design principle is 'checker, not a maker': the system supports the lawyer's judgement and does not send AI-generated work product directly to clients. The proposal also contemplates client-specific data segregation and other security controls.

2

Misfits Margin

Issue addressed: AI-assisted drafting can accelerate document production, but it may also reduce the iterative drafting and review through which junior lawyers traditionally learn professional judgement. Supervisors' edits often contain valuable reasoning, yet that reasoning can remain tacit or disappear once a document is finalised.

Proposed approach: Margin is an AI-enabled learning platform that captures the learning embedded in supervisors' edits and feedback. It compares a junior lawyer's draft with the supervisor's markup, identifies substantive changes and annotations, and converts them into structured explanations and reusable learning resources that help the junior understand not only what changed but why. The tool is intended to fit into existing review workflows rather than create a separate training exercise, allowing firms to preserve supervisory knowledge and apprenticeship while still benefiting from faster AI-assisted drafting. Human judgement remains central, and the proposal highlights the need to protect confidentiality and avoid turning the tool into a performance-monitoring mechanism.

3

Chubbles CiteCheck SG

Issue addressed: Generative AI can produce fabricated, inaccurate or outdated legal citations, while lawyers remain professionally responsible for authorities cited in their work. Manually checking every citation can erode the productivity gains that AI-assisted drafting is intended to provide.

Proposed approach: CiteCheck SG is a Singapore-focused citation verification platform. It extracts citations from draft legal documents and checks them against a structured citation graph built from Singapore judgments, returning transparent outcomes such as Confirmed, Flagged, Suspect or Uncertain. Rather than relying on generative inference to decide whether an authority is valid, the verification process uses deterministic and graph-based evidence so that the result can be traced and reviewed. The proposal is designed to sit within existing drafting workflows, reducing manual verification effort while preserving the lawyer's responsibility to review and confirm authorities before use.

4

Team Crank From Principles to Practice

Issue addressed: Singapore already has high-level guidance on responsible AI use, but firms - particularly smaller practices - may lack the time, expertise or implementation capacity to translate principles into everyday governance, pricing, training and learning practices.

Proposed approach: From Principles to Practice is a modular implementation package that helps firms operationalise responsible AI rather than create another set of principles. It brings together four connected components: an AI governance framework covering approved uses, safeguards and accountability; a fee-model framework addressing how AI-enabled efficiencies can be reflected in charging arrangements; an AI capability pathway for developing staff skills; and a prompt experiment log with periodic organisational learning reviews so that experience with AI is captured and shared. Firms could adopt the components progressively according to their maturity and circumstances, using existing Singapore guidance as the foundation.

5

Jam & Raph Form Filler

Issue addressed: Preparing court forms can involve substantial mechanical work: extracting information from source documents, entering the same facts across large numbers of fields, checking them manually and re-keying information into filing systems. This creates avoidable administrative burden and opportunities for input errors.

Proposed approach: Form Filler is designed to draft Singapore court filings from the documents lawyers already hold, beginning with the First Ancillary Affidavit. The lawyer uploads the relevant source documents; the system extracts information for the required fields and prepares the form, with each value available for review before submission. The proposal deliberately keeps the lawyer in control: the workflow is designed to require human confirmation of the filing, limit use of client documents to the filing purpose, and provide visibility over extracted information. Its broader aim is to automate repetitive form-filling without requiring lawyers to learn complex prompting or surrender oversight of client data and final submissions.

6

WongP Clawset

Issue addressed: Junior lawyers can spend significant time locating, organising and summarising matter information instead of performing substantive legal work. Using more autonomous AI systems for these tasks may create additional concerns around accuracy, source verification and unintended actions.

Proposed approach: Clawset is a read-only AI matter assistant focused on information corralling rather than autonomous legal analysis. It retrieves information from matter repositories through natural-language queries, prepares source-linked matter briefings and timelines, tracks deadlines and supports knowledge management through precedent updates. Every output is tied back to the underlying source material through pinpoint citations so that lawyers can verify the information quickly. By remaining read-only and avoiding autonomous changes to documents or systems, Clawset seeks to capture productivity gains from AI-assisted information management while maintaining transparency and professional oversight.

7

Red Pulse Agent Red

Issue addressed: Time recording remains a frequent, manual and fragmented administrative task. Lawyers may reconstruct their day from memory, which can lead to missed time, inconsistent narratives and unnecessary administrative effort, while firms and clients need accurate and defensible billing records.

Proposed approach: Agent Red is a Copilot-based time-recording agent that reconstructs a lawyer's workday from Microsoft 365 activity such as emails, documents, meetings and workspace context. Using firm-approved prompts, templates, access rules and output standards, it groups activity into tasks, estimates time and generates a structured draft daily time record with billing narratives. The lawyer reviews and edits the draft before approving output in a format compatible with the firm's billing system. The proposal deliberately targets a lower-risk administrative use case: it does not provide legal advice or drafting, relies on review before submission and uses traceable activity metadata to support the resulting entries.

8

WP-KIT RADAR

Issue addressed: Lawyers can miss business opportunities, legal and regulatory developments or important action items because relevant information is spread across news sources, legal sources, email and documents. At the same time, confidentiality concerns can discourage use of AI for high-value workflows.

Proposed approach: RADAR is an ambient AI layer designed to work inside applications lawyers already use, surfacing relevant information and suggested actions without requiring repeated prompting or context switching. It comprises four scans: Opportunity Radar monitors curated news and deal sources for potential client opportunities; Knowledge Radar tracks legislation, case law and regulatory developments; Action Radar classifies incoming messages and surfaces prioritised tasks; and Trust Shield screens documents and communications for personal or sensitive information. Trust Shield is designed for local processing so content need not be transmitted to an external server. The system adapts outputs to different roles and provides one-click follow-up actions while leaving the user in control.

9

TSMP Disputes AI Twinning

Issue addressed: Lawyers have limited opportunities to rehearse advocacy, test arguments and obtain preliminary feedback before appearing in court or seeking review from senior practitioners. Junior lawyers may also hesitate to ask basic or repeated questions when supervisory time is constrained.

Proposed approach: AI Twinning proposes digital simulations of judges, senior practitioners and other legal actors built from publicly available or appropriately authorised written materials. Lawyers could use these simulations to test arguments, anticipate questions, refine submissions and practise advocacy before human review. The proposal also envisages internal firm twins that provide junior lawyers with accessible preliminary guidance before work is escalated to supervisors. The twins are framed as simulations rather than substitutes for professional judgement, with human oversight and governance safeguards required around matters such as consent, data use, validation and the boundaries of judicial profiling.

10

NextGen Legal LawLock

Issue addressed: AI-assisted legal drafting creates a risk that statutory references may be inaccurate, outdated or contextually wrong. Lawyers can respond by manually checking every reference, but doing so reduces the efficiency gains of AI and may be especially burdensome for smaller practices and in-house teams.

Proposed approach: LawLock is a Singapore-focused citation verification engine that audits AI-generated legal drafts against Singapore legal sources. Its prototype extracts statutory references using NLP and pattern matching, applies deterministic and semantic matching, flags invalid or questionable references and presents verified citations or suggested corrections for lawyer review. A feedback mechanism captures lawyer-approved corrections, which are validated before they update matching rules, allowing the system to improve over time without bypassing professional review. The proposal is positioned as a specialised safeguard for Singapore legal work rather than a general-purpose drafting system.

11

Condition Precedent Docklet: The First Matters Companion

Issue addressed: First consultations can be inefficient when clients arrive with unstructured narratives, incomplete information and large sets of documents. Lawyers may spend significant time during or before the consultation identifying what is relevant, while clients may not know what information to prepare.

Proposed approach: Docklet is an AI-powered pre-consultation companion that helps structure information before the first substantive meeting. It guides clients through matter-specific intake, organises uploaded information and documents, and uses firm-curated playbooks and retrieval-augmented generation to identify relevant issues, gaps and follow-up questions. The resulting structured material is intended to prepare both client and lawyer for a more focused consultation, rather than replace the lawyer's advice. By anchoring the workflow in firm-curated knowledge and keeping the substantive consultation with counsel, the proposal seeks to improve intake efficiency while maintaining professional oversight.

12

USP Law Offline Redaction Tool

Issue addressed: Lawyers may wish to use external generative AI tools but remain concerned that confidential or sensitive client information could be exposed before or during processing. Manual redaction can be time-consuming and may remove much of the productivity benefit of using AI.

Proposed approach: USP Law proposes an offline redaction and substitution tool that operates locally before any material is sent to an external AI model. The programme identifies defined categories of confidential information, replaces them with neutral placeholders and displays the sanitised prompt for lawyer review. Only after the lawyer is satisfied is the redacted material submitted to the chosen AI model; the returned response can then be mapped back to the original context for review. The tool is deliberately separate from the external AI platform, with local processing and no plugin connection before redaction. The proposal also recognises the risks of under- or over-redaction and allows adjustable redaction depth and manual override.

13

OptiNUS Prime Raising the Bar with Artificial Intelligence

Issue addressed: AI adoption support already exists through grants, implementation assistance and ethical guidance, but AI literacy can still depend heavily on individual motivation and the resources of each firm. This may create uneven capability across the profession, particularly between larger firms and small and medium-sized practices.

Proposed approach: OptiNUS Prime proposes treating AI literacy as a profession-wide competency rather than an optional skill. Its core proposal is a mandatory AI literacy module integrated into the Part B course for new entrants, complemented by ring-fenced AI-related continuing professional development for existing practitioners. The curriculum would give lawyers a common baseline in effective AI use, limitations, verification, confidentiality, professional responsibility and human oversight. By embedding learning within existing professional education structures, the proposal aims to create a baseline that does not depend on firm size or access to internal technology teams.

14

ANGUS Prompt-Building Tool

Issue addressed: Legal professionals may recognise the potential of generative AI but struggle to translate a legal task into a sufficiently clear, structured and detailed prompt. Repeatedly constructing effective prompts can also be cumbersome, creating a practical barrier to consistent use.

Proposed approach: ANGUS proposes a guided prompt-building tool that helps lawyers construct stronger prompts without requiring them to master prompt engineering from scratch. The interface breaks a request into structured components and guides the user through relevant choices and inputs before assembling them into a more complete prompt for a generative AI system. The approach is intended to make prompting more repeatable and accessible, reduce the effort involved in starting from a blank text box and encourage users to provide the context and instructions needed for more useful outputs. The tool assists with prompt construction rather than replacing the lawyer's responsibility to assess the resulting AI output.