



Source of Wealth Identification: Agentic AI for Source of Wealth Due Diligence in Private Banking

RANASINGHE, Yehan Lakvindu
Advised by Prof. XIAO Huiyu



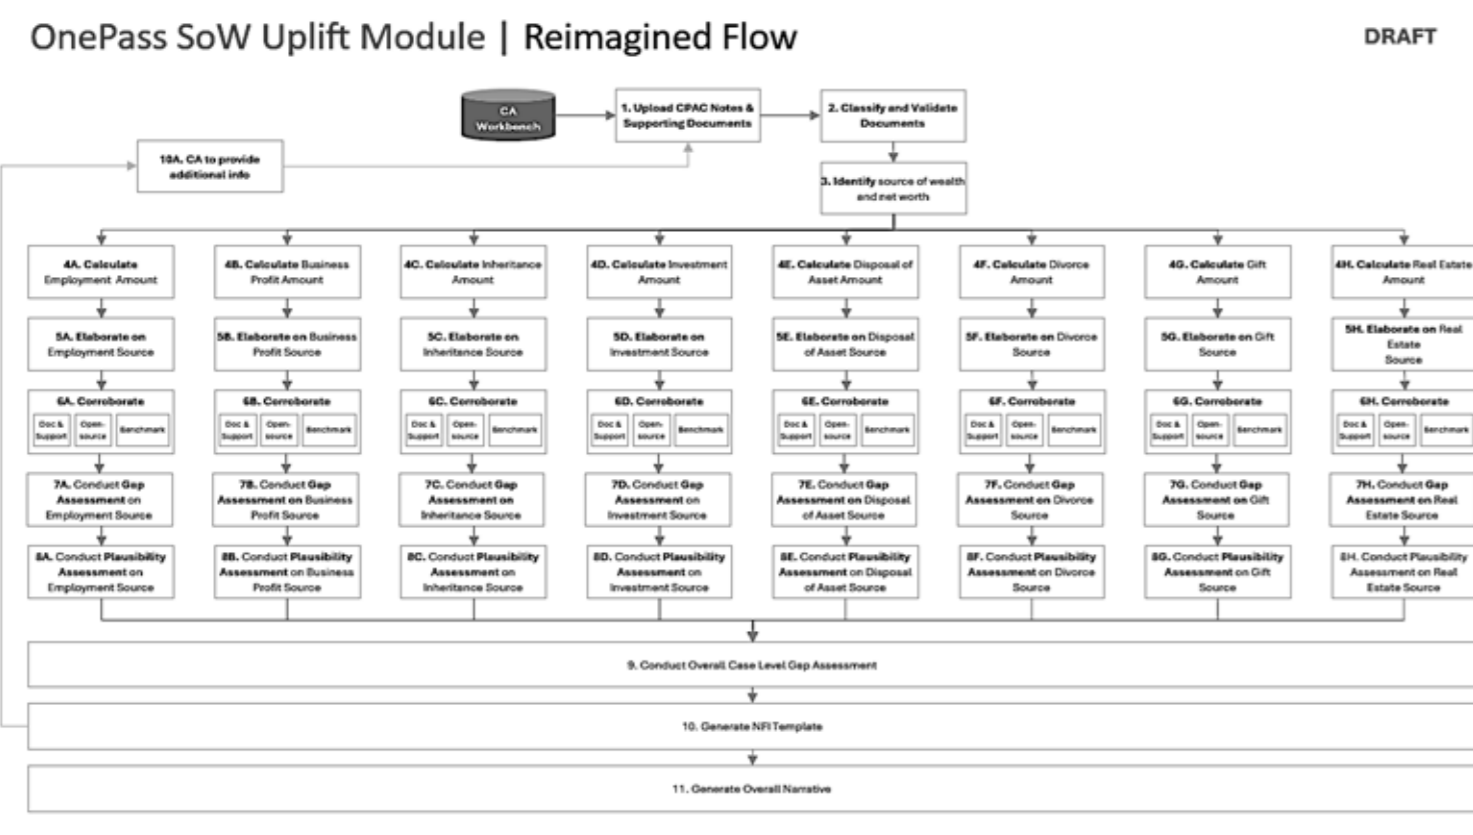
What is SoW Due Diligence?

Source of Wealth (SoW) due diligence is a critical component of customer onboarding and ongoing monitoring in private banking and wealth management. However, traditional SoW reviews are highly manual, resource-intensive, and subject to reviewer inconsistency. This project proposes an Agentic AI-based SoW analysis platform built by Deloitte for their clients.

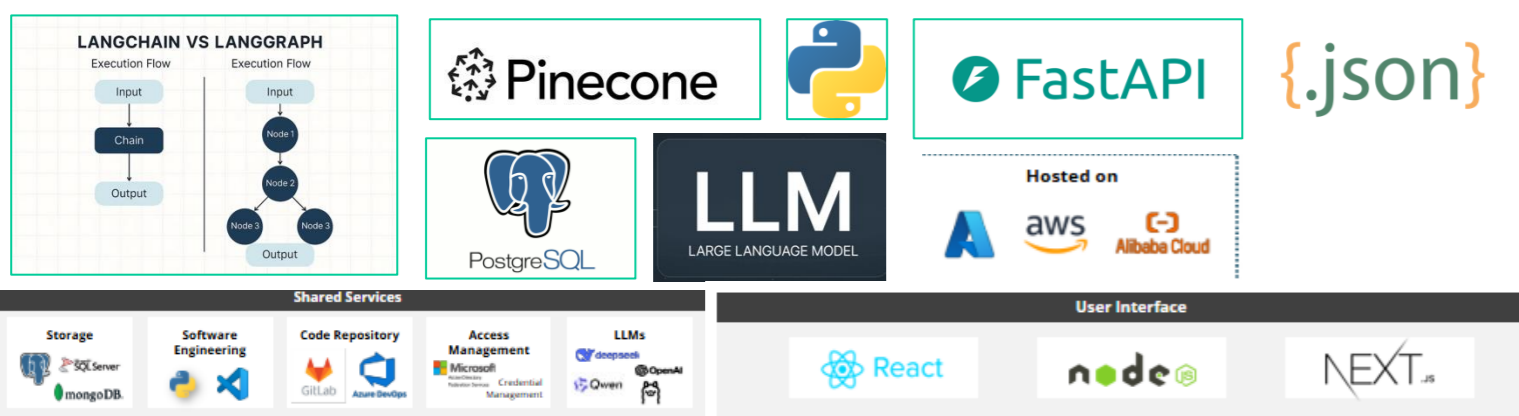
Multi-Agent SoW Analysis System

A modular agent pipeline designed to simulate expert SoW review tasks:

- **Identification Agent** – extracts wealth categories
 - **Elaboration Agent** – expands and structures narratives
 - **Corroboration Agent** – cross-checks claims against documents
 - **Gap Analysis Agent** – identifies missing or inconsistent evidence
 - **Plausibility Agent** – assesses overall credibility and risk
- Each agent produces structured, auditable outputs.



LangGraph is used to orchestrate agents as nodes in a directed graph, where each agent performs a well-defined task and passes structured outputs to subsequent agents. This design closely mirrors how human compliance teams perform step-by-step SoW reviews.



1. Road Runner Case Scoring & Routing Mechanism

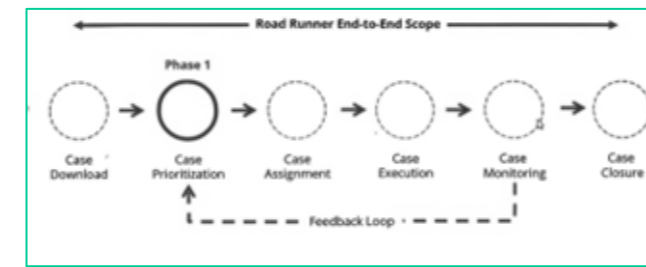
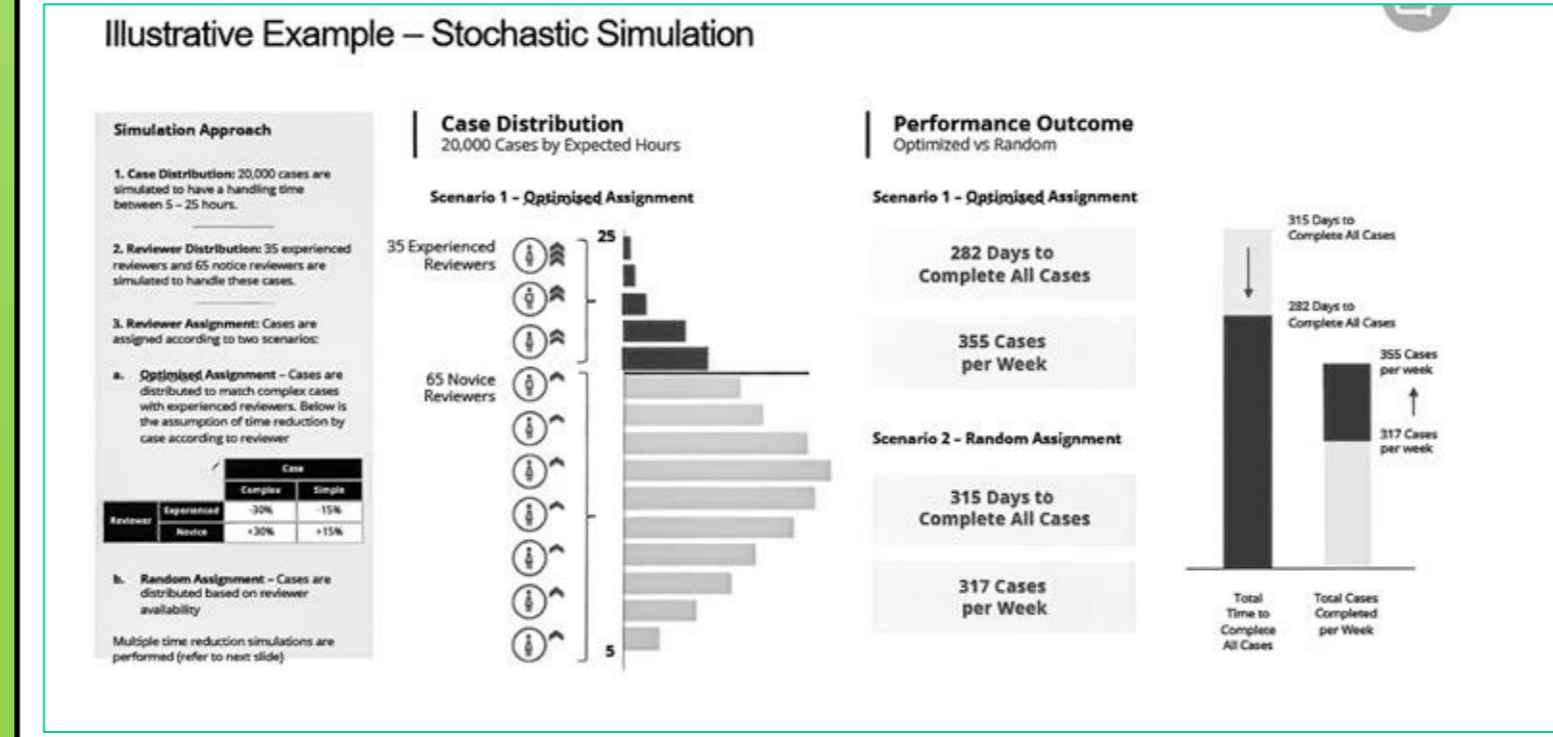
Road Runner translates agent outputs into operational decisions: It Aggregates outputs from all agents, Assigns a complexity score to each SoW case, and then routes cases to either junior reviewers (low complexity) or Specialist reviewers (high complexity). This optimizes reviewer workload and reduces bottlenecks.

How Road Runner Works

- Aggregates outputs from multiple agents
- Computes a composite complexity and risk score per case
- Low complexity → junior or expedited review
- High complexity → senior reviewer or escalation queue

Key Benefits

- Optimizes reviewer workload allocation
 - Reduces bottlenecks in compliance operations
 - Ensures complex cases receive scrutiny
- Road Runner acts as a decision layer

2. Back testing Framework Overview

The back testing task verifies whether each claim or statement generated by the AI can be directly traced to supporting evidence within the client-provided documents. This ensures factual grounding, minimizes hallucinations, and supports regulatory audit requirements.

Claim-to-Document Mapping

- Each AI-generated output is decomposed into atomic claims
- For each claim, the agent searches across the ingested client documents
- The agent identifies:
 - The source document
 - The exact text span or paragraph where the claim is supported

Role in Evaluation

- The back testing agent feeds its input into the Proxy Question (PQ) framework
- PQs assess whether:
 - Every material claim is evidence-backed
 - Sources are correctly cited including the page number
 - If a claim is unverified, then the agent flags them as Not Verified

Wealth Corroboration Report

