

AI Co-Pilot for High-Net-Worth Underwriting

Customer:

A Leading Pan-Asian Life Insurer

Decisioning Capability

Autonomous Case Planning & Disposition

What it does:

From a single goal, the Co-Pilot plans and executes the entire HNW underwriting flow without step-by-step instructions.

How it works:

- Ingests and summarises unstructured inputs (application, APS, labs, financials, AML/PEP).
- Retrieves guideline evidence via RAG with versioned citations (impairment manual, country capacity, treaties, AML policy).
- Scores risks (medical/financial/AML) and checks capacity & facultative triggers.
- Detects evidence gaps against age/amount grids and proposes precise next steps.
- Recommends disposition (Accept / Table / Refer / Decline) with audit-ready details.

Assurance targets:

- Plan completeness of canonical steps
- Decision consistency vs senior underwriter gold labels
- 100% of recommendations fully cited

Actioning Capability

Tool Selection, Use & Safe Execution

What it does:

Selects and orchestrates the right tool for each sub-task, switches when complete, and executes safely within policy.

How it works:

- Application APIs for policy/member/product rules
- Connector toolkit for ecosystem integrations (payments, email, comms)
- Data fabric lookups for policy and eligibility checks
- Rules and decision table execution for business logic
- Search/KMS for PDF or page retrieval
- ML models for ICD inference, document classification, and data extraction
- LLM (RAG-bound) for grounded answers & NBAs
- Code generation & execution for capacity/treaty math
- Versioned knowledge DB for retrieved snippets

How it behaves:

- Intake & classification: Classifier labels each file by type and language.
- Evidence extraction with privacy: PHI redactor masks identifiers before LLM processing; originals stored securely.
- Grounded retrieval & reasoning: Queries versioned KB and composes NBAs strictly from retrieved content.
- Enterprise actioning & math: Runs eligibility checks, treaty math, and creates tasks (e.g., order Stress Echo, request APS addendum).

Perceiving Capability

Environment-Aware Re-planning

What it does:

Observes new information and adapts the plan mid-flight without human prompting.

Perception & responses:

- New lab (low HGB): Re-score risk, add anemia workup, refresh disposition.
- Name/DOB mismatch: Pause, trigger NIGO outreach, resume on correction.
- Guideline update: Re-retrieve, re-score, escalate if thresholds change.
- API outage: Switch to cached embeddings, flag citations as cached.
- Contradictory APS note: Raise evidence-conflict task with highlighted spans.

Assurance targets:

- Correct re-plan in most runs
- Improved output quality after re-plan
- Tracked decision deltas post-adapt

