

● **HALLEY · STRUCTURED. GOVERNED. REVIEWABLE.**

Halley. Trustworthy AI for actuarial work.

Most AI tools can generate an answer. The real question is whether you can trust how it was produced. **Halley is the StrinGaze Actuarial AI Agent — real software running today**, working through structured models and governed knowledge so teams move faster without giving up control.

Four agent capabilities are live in the platform now. Each keeps the actuary in the loop — Halley drafts, suggests, and explains; you accept, reject, or refine.

Model Analysis

Ask about variables, tables, formulas, and structure; trace how inputs flow through the calculation chain

Help Manual

Step-by-step answers grounded in StrinGaze product guides and manuals — not generic web text

Model Checkup

Run quality rules on a selected model for a health score, severity ranking, and fixes to review

Smart Modeling

Plan, then apply model changes in your workspace — plan → confirm → execute, with the actuary in control

THE FOUR LIVE AGENTS — TODAY

Halley

The AI agent layer



Doer agent

Executes the modelling or analysis task



Checker agent

Reviews logic, structure, and output before acceptance



Best-practice recipes

Embedded workflow patterns for consistent execution



Governed validation

Checks results against structure, rules, and expected patterns

API / MCP

How Halley reads & operates your models



Input API

Read assumptions, tables, model points



Output API

Surface results back to the actuary



Model Management

Versioned, branchable, auditable



Task Submit

Trigger runs, track progress

StrinGaze Foundation

The structural primitives Halley reasons over



Class Tree

AI-native inheritance — ~80% less table data & code, easy for AI to interpret



Smart Table

Structured assumptions and table data, separated from logic for cleaner AI reasoning



Knowledge base

Governed prompts, rules, and mappings to keep outputs more consistent across LLM choices



Governed workflow

Reviewable orchestration designed to keep modelling work simple, transparent, and repeatable

Running today. Expanding fast.

The last decade was about producing trusted numbers. The next is about turning actuarial capability into **enterprise intelligence**. Halley starts today as an assistant that answers questions — and grows, phase by phase, into an agent layer that helps **orchestrate decisions**. Every phase keeps the actuary in control.

TODAY

Current state — running software

What Halley does today.

- **Model Analysis** — explains variables, tables, formulas and Class Tree structure; traces how inputs flow through the calculation chain.
- **Help Manual & Model Checkup** — step-by-step answers grounded in product docs; quality rules return a health score and ranked fixes.
- **Smart Modeling** — plans, then applies model changes in your workspace: plan → confirm → execute, actuary in control at each step.

WHAT'S COMING

From answering questions → orchestrating decisions

Where Halley is going.

- **Phase 1 · Design-to-Model** — design documents become runnable models: graded structured parameters, mapped to a master product, run on a dual basis. *The foundation.*
- **Phase 2 · Pricing & FP&A agent layers** — agents that orchestrate analysis, planning and steering across the lifecycle. *Acceleration.*
- **Phase 3 · Governed Decision Intelligence** — the Governance Agent and enterprise foresight turn outputs into trusted, forward-looking decisions. *Leadership.*

HOW HALLEY STAYS GOVERNED

DOER + CHECKER

Our AI agent is split into a doer and a checker. One executes, the other reviews logic, structure and result — and the actuary is always the final reviewer.

AI-NATIVE CLASS TREE

~80% less table data and code. Our Class Tree structure is AI-native and easy for AI to understand — it reasons about the model, not just the text.

GOVERNED, LLM-AGNOSTIC

Our knowledge base is built on governance to ensure consistent results no matter which LLM you run — and you can swap models freely.

BEST-PRACTICE BY DESIGN

Embedded best-practice structure means anyone who builds gets the same result — models stay simple, transparent, well-maintained, and far less prone to hallucination.

GOVERNANCE AGENT

An independent gatekeeper (roadmap) enforces source validity, red-lines, required confirmations, and traceability before release.

YOUR DATA STAYS YOURS

Halley runs where StrinGaze runs — public cloud, private cloud, or on-prem. Your data never trains anyone else's model.

WHY HALLEY IS DIFFERENT

Can you really trust AI built on an unstructured model, without governance?

Other tools claim they can do this and that. But AI on a flat, unstructured model — with no governance — can read text, not reason about it. **Halley is built on the structured Class Tree, governed knowledge, and a doer-checker agent embedded with best-practice structure**, so models stay simple and transparent. Non-standard models, by contrast, accumulate inefficiency and need periodic optimisation exercises just to stay workable.

StrinGaze Pte. Ltd.

contacts@stringaze.com · linkedin.com/company/stringaze-sg · www.stringaze.com

* Capabilities described under "Today" are demoable now; "What's coming" items are subject to development and pilot validation.



READ ONLINE