AI agents for the work around the rule and the control

Agents that read your policies, your registers, your control evidence and your incident records, prepare the work, and stop for a named officer. On infrastructure you control.

THE PROBLEM

Of 15 public risk and compliance workflow templates we opened, 2 hold every action for a person, 5 do so in part and 8 not at all. Of 111 public interface descriptions, 103 import and 8 are over the import limit; the 103 give 16,041 operations, 8,667 sent as changes, so tools are chosen one by one.

HOW IT WORKS, IN THREE STEPS

STEP 01

Connect what the agent may read.

Policies, procedures, the control library, the obligation register and past audit reports become knowledge sources, with each reader's permissions kept. The risk and control system, the incident log and the case system are reached through the description of an interface you already run; rules, evidence and questionnaires are uploaded by the officer who holds them.

STEP 02

Decide where it stops.

Every tool call passes your policy. A tool that would record a rating or a test result, close a finding or file anything waits for a named person.

STEP 03

Run it where the registers live.

In your data centre or your private cloud, on a model you serve.

THE RULE EVERY AGENT FOLLOWS

It prepares. A named risk or compliance officer decides.

Thirty-five agents in nine areas

Regulatory change	Regulatory Monitoring Agent. Regulatory Change Impact Note. Obligation Register Draft. Consultation Response Draft
Policies	Policy Q&A. Policy Gap Check. Policy Update Draft. Attestation Follow-up
Risk register and assessments	Risk Register Check. Risk Assessment Preparation. Risk Indicator Report. Risk Event Summary
Controls and testing	Control Test Preparation. Control Evidence Check. Control Mapping. Finding Follow-up
Incidents and notifications	Incident Log Draft. Notification Deadline Check. Regulatory Notification Draft. Breach Register Summary
Third parties	Third-Party Assessment Preparation. Third-Party Register Check. Questionnaire Answer Draft
Conduct and cases	Case Intake Summary. Investigation File Summary. Declarations Register Check. Training Record Check
Audit	Audit Planning Pack. Evidence Request List. Working Paper Draft. Audit Finding Draft
Reporting	Regulatory Return Preparation. Committee Report Draft. Compliance Plan Status. Document Generation

Also in the pack: the full list of what each agent reads and where it stops, and the white paper “AI agents inside the risk and compliance department”, with sources.

Seven questions to ask before an agent touches a risk register, a control test or a regulatory filing

01 Where does it run, and where is the model?

02 Does it act as the person using it, with that person’s permissions?

03 Which tools record a rating or a test result, close a finding or file anything, and does each wait for a person?

04 Are the read operations and the change operations separate tools?

05 Does it pass or fail a control or decide what is notified, or does an officer?

06 Does it say what it could not read, and when it found nothing?

07 What is recorded, and can we stop one agent or all of them at once?

WHERE THE PLATFORM RUNS TODAY

In production at Taiwan Cooperative Bank, air-gapped, inside the bank’s own perimeter.

MODON built its AI Centre of Excellence on Katonic inside the Kingdom of Saudi Arabia.

ePLDT, the ICT arm of the Philippines’ largest telco, selected Katonic to power Pilipinas AI, a sovereign AI cloud hosted in the Philippines.

Where to start

One team, one rule or one control, invented data first. A policy question and answer agent on your policies can run on the first day, because no register is touched.

KATONIC.AI/SOLUTIONS/RISK-COMPLIANCE