AI agents for the work around the alert and the case

Agents that read your alerts, your logs, your scan results and your playbooks, prepare the work, and stop for a named analyst. On infrastructure you control.

THE PROBLEM

Of 14 public security workflow templates and playbooks we opened, 4 put an analyst's approval before a containment action, 3 do so in part and 7 not at all. Of 107 public interface descriptions, 89 import, 16 are over the limit and 2 are refused; the 89 give 15,450 operations, 7,540 sent as changes.

HOW IT WORKS, IN THREE STEPS

STEP 01

Connect what the agent may read.

Playbooks, security policies, the asset list, past incident reviews and the detection rules become knowledge sources, with each reader's permissions kept. The alert and log platform, the case system, the scanner and the directory are reached through the description of an interface you already run; alert exports, reported emails and scan reports are uploaded by the analyst who holds them.

STEP 02

Decide where it stops.

Every tool call passes your policy. A tool that would isolate a host, disable an account, block an address or close an alert waits for a named person.

STEP 03

Run it where the logs live.

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

THE RULE EVERY AGENT FOLLOWS

It prepares. A named analyst decides.

Thirty-five agents in nine areas

Alerts	Alert Triage Summary. Indicator Lookup. Related Alert Check. Shift Handover Summary
Reported emails	Reported Email Triage. Similar Email List. Reporter Reply Draft. Reported Email Summary
Incidents	Incident Response. Incident Timeline. Incident Record Draft. Post-Incident Review Draft
Vulnerabilities	Vulnerability Prioritisation Preparation. Advisory Impact. Patch Exception Check. Scan Result Summary
Identity	Sign-in Investigation Summary. Privileged Access Review Preparation. Dormant Account List. Access Change Summary
Cloud and configuration	Configuration Finding Summary. Baseline Comparison. Exposed Resource List
Threat intelligence and detection	Threat Report Summary. Detection Coverage Check. Detection Rule Draft. Rule Tuning Report
Assurance	Security Questionnaire Answer Draft. Test Finding Follow-up. Security Exception Register Check. Security Audit Evidence
Awareness and reporting	Knowledge. Awareness Note Draft. Security Report. Notification Deadline Check

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

Seven questions to ask before an agent touches an alert, an account or a detection rule

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 isolate a host, disable an account, block an address or close an alert, and does each wait for a person?

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

05 Does it decide what an alert or a reported email is, or does an analyst?

06 What it reads may be written by an attacker. Which tools could act on that text with no person?

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 kind of alert, invented data first. A knowledge agent on your playbooks can run on the first day, because no system is touched.

KATONIC.AI/SOLUTIONS/SECURITY