AI agents for the work around the draft and the claim

Agents that read your brand guide, your approved messaging and your product facts, draft and check the work, and stop for a named person. On infrastructure you control.

THE PROBLEM

A regulator has said the advertiser answers for what its tools produce. In the vendor documents we read, one has a person approve every post, and one has the agent publish a campaign once its plan is approved. Of 145 public interface descriptions from 26 publishers, 142 import and 3 are too large (31 more were left out first: wrong format). The 142 carry 5,554 operations, 2,942 sent as changes.

HOW IT WORKS, IN THREE STEPS

STEP 01

Connect what the agent may read.

The brand and style guide, approved messaging, product facts, the claims register and past content become knowledge sources, with each reader's permissions kept. Pages, campaigns, assets and events are reached through the description of an interface you already run; results through a reporting database, read only.

STEP 02

Decide where it stops.

Every tool call passes your policy. A tool that changes something waits for a named person. No agent is given a tool that publishes a page, posts, schedules or sends a campaign, or changes an advert or a budget. Research is about companies.

STEP 03

Run it where the data lives.

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

THE RULE EVERY AGENT FOLLOWS

It prepares. A named person decides.

Thirty-five agents in nine areas

Content	Content Brief. Content Drafting. Brand and Style Check. Channel Version
Website and search	Search Visibility Report. Page Title and Description. Site Content Audit. Landing Page Drafting
Social media	Social Post Drafting. Social Plan. Comment Summary. Comment Reply Drafting
Email and campaigns	Campaign Brief. Email Drafting. Campaign Check. Campaign Results Report
Events and webinars	Event Plan. Speaker Brief. Event Follow-up Drafting. Session Summary
Product marketing	Launch Plan. Messaging Check. Customer Story Drafting
Claims, assets and approvals	Claim Check. Testimonial Check. Asset Label Check. Approval Pack
Market and accounts	Market Research Summary. Competitor Messaging Summary. Account List. Internal Newsletter
Operations and reporting	Marketing Report. Budget Report. Translation Draft. Lead Handover 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 marketing team”, with sources.

Seven questions to ask before an agent drafts for the public or reads a campaign

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

02 Does it keep each person’s permissions?

03 Which tools change something, and does each wait for a person?

04 Can any agent publish, post, schedule or send without a person?

05 Where does each claim come from, and who labels AI-made content?

06 Does it profile a named person or score a marketer?

07 Can we stop one agent, one tool, or all of them, at once?

WHERE THE PLATFORM RUNS TODAY

In production at a top-tier bank in Asia, 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 channel, invented data first. A brand and style check can run on the first day: it touches no system.

KATONIC.AI/SOLUTIONS/MARKETING