zango

How to build AI agents for compliance: marketing review and approval

A masterclass with Ninety One and the International Compliance Association

01 Introduction

02 The live build

03 In conversation with Ninety One

04 Q&A and sandbox access

About us



Ashi Bajwa

Regulatory AI Architect

zango



Sam Green

Senior Policy & Partnerships Lead

zango



Matthew Huntington

Head of Marketing and Distribution Compliance

Ninety One



James Booth

Principal Advisor

ICA

zango

POLL

Where does AI sit in your marketing review today?

The problem with marketing review and approval



Handoffs between marketing and compliance



Keeping up with change



Auditability



Accuracy of general purpose tools

zango

The anatomy of a good AI use case



High-volume / repetitive



Clear source of truth



Defined output structure

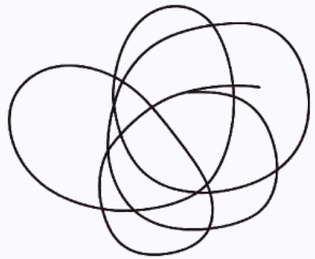


Auditable

Financial promotions review as an AI use case

- High volume and repetitive
- Clear source of truth
- Defined output structure
- Auditable

Let’s build a financial promotions review agent



The use case: Fund marketing



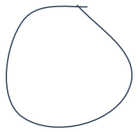
A professional fund marketing asser

A UK asset manager is refreshing the marketing for a professional fund: needing sign-off.



Review is the bottleneck

Every asset goes back and forth between marketing and compliance, with the same checks repeated on each new version.



A first-pass review agent

The agent reads each asset, checks it against the rules and approved claims, and drafts fixes for a person to approve.

The agent, step by step

01 Context ingestion

Reads the whole asset, including text, visuals and audio, and identifies the product it promotes.

02 Rules review

Finds the regulatory and internal rules in scope and flags where the promotion falls short.

03 Claims library

Matches each claim to approved and rejected claims, with evidence, caveats and expiry, and flags new ones.

04 Suggested fixes

Drafts amends in your firm's language: risk warnings from past decisions, pre-approved wording and repeat corrections.

Where the human steps in

Compliance reviews the flags

Checks regulatory and internal rule breaches, and any previously rejected claims.

Marketing revises the asset

Makes changes based on the findings and suggested fixes, before it goes back for review.

The claims library stays current

Newly approved claims go into the library, so the next review starts further ahead.

Compliance owns final sign-off

The agent drafts and flags. A person approves every promotion.

zango

POLL

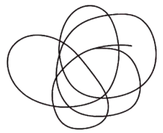
Which stage of marketing review takes your team the most time?

Limitations to be aware of



Maintenance

Rules, claims and products change. A self-build needs someone to keep the rulesets and claims library current.



Accuracy

General purpose tools can give different answers to the same asset. Review needs the same result every time.



Auditability

General purpose AI tools aren't designed to manage compliance approval workflows.

zango

INDUSTRY INSIGHTS

In conversation with Ninety One

What is changing in fund marketing

01 **The new CCI product summary**

Mandatory for every UK retail fund by 8 June 2027, with no template and no exemption from the financial promotion rules.

02 **Approval and consumer outcomes**

Financial promotion requirements, sign-off responsibilities, and testing documents for different investor profiles.

03 **AI in marketing compliance**

First-pass reviews, consistency checks across fund documents, and practical applications of AI agents.

04 **Building the compliance workflow**

Build versus buy, integrating AI into existing processes, and lessons from Ninety One's experience.

zango

POLL

Would you build or buy a marketing review agent?

zango

Questions



zango

Try it on your own assets



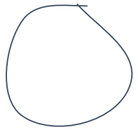
48 hours in the sandbox

Every attendee gets 48 hours after the session to test the workflow on their own marketing.



Work emails only

Sign up with your work email. Access is limited to work domains.



Up to two assets

Upload two of your own assets and see the full review, from rules check to suggested fixes.

The sandbox runs a sample set of UK investment marketing rules (COBS 4, GEN 4 and 5, ESG 4). It isn’t exhaustive, doesn’t include your internal rules or claims library, and its findings are a first pass, not a sign-off. To test the platform’s full capabilities, get in touch.

Contact us



Ashi Bajwa

Regulatory AI Architect, Zango

ashi@zango.ai



Sam Green

Senior Policy & Partnerships Lead, Zango

sam@zango.ai

zango

RESOURCES

Agentic build prompts

Using the build prompts

Copy a prompt into Claude or ChatGPT and replace each **{{ highlighted slot }}** with your own material. Run them in order: Context first, then the three checks (they can run at the same time), then Suggested fixes with the findings from the checks. Run Review report last, once a person has decided on each finding. In the canvas the promotion lines come from reading the text off the ad image; outside it, type the ad's text one line per row.

01

Context

Reads the promotion alone, with no rules. Works out what it is and pulls out every claim, so every check that follows looks at the same list.

Reads: The promotion

02

Regulatory review

Reads the promotion, the claims found and the FCA rules wired into it. Every finding must quote the promotion and cite one of those rules.

Reads: Promotion, Context claims, FCA rules

03

Internal rules review

Same mechanics as the regulatory review, against the firm's own rulebook. This is where house style and past lessons live.

Reads: Promotion, Context claims, internal rules

04

Claims library check

Checks each claim against the claims library: approved, expired, rejected before, or not in the library at all.

Reads: Promotion, Context claims, claims library

05

Suggested fixes

Grounds each fix in approved wording, approved claims and past decisions. It drafts new wording only when nothing approved fits, and says so.

Reads: All findings, approved wording and past decisions

06

Review report

Runs after human review. Writes up the findings, fixes and the reviewer's decisions in the standard review report format. It reports what was decided and does not review the promotion again.

Reads: Context output, all findings, fixes, reviewer decisions

Agent 1: Context

Reads: The promotion

Reads the promotion alone, with no rules. Works out what it is and pulls out every claim, so every check that follows looks at the same list.

You are helping a UK financial promotions compliance reviewer. Work only from the material below. Do not invent rules, claims, IDs or facts.

TASK
Read the promotion. Work out the product, the likely audience and the channel. List only the marketing claims: statements about returns, risk or safety, quality, awards, sustainability, cost or ease. Quote the exact words. Leave out the brand name and the regulatory status line.

PROMOTION (text read from the image, line by line, top to bottom. size = text height as a % of the image height, so bigger means more prominent. top = how far down the image the line starts. The reader can garble very large or stylised text, so read odd characters with care.)
{{ one line per text line of the ad: L1 | text | size 6.7% | top 21% }}

CAPTION OR NOTES (only if the post has a caption)
{{ caption text }}

OUTPUT
Reply with only JSON in this shape:
{"product":"...","audience":"...","channel":"...","claims":[{"quote":"exact words","line":"L6"}]}
Copy each quote exactly as it appears in the promotion, word for word. Keep quotes short: the claim itself, not the whole sentence.
For every quote, give "line": the ID of the text line it comes from, for example "L5". If a quote runs over two lines, give the first.

Agent 2: Regulatory review

Reads: Promotion, Context claims, FCA rules

Reads the promotion, the claims found and the FCA rules wired into it. Every finding must quote the promotion and cite one of those rules.

You are helping a UK financial promotions compliance reviewer. Work only from the material below. Do not invent rules, claims, IDs or facts.

TASK
Check the promotion against the rules below. Flag only clear breaches that a compliance reviewer would ask marketing to fix. Quote the shortest exact words that cause each breach and give the single most specific rule. If the same words break several rules, report them once. Use COBS 4.2.1R only when no more specific rule fits. Something required but missing, such as a risk warning, is one finding with the quote "missing". At most 4 findings, most serious first.

PROMOTION (text read from the image, line by line, top to bottom. size = text height as a % of the image height, so bigger means more prominent. top = how far down the image the line starts. The reader can garble very large or stylised text, so read odd characters with care.)
{{ one line per text line of the ad: L1 | text | size 6.7% | top 21% }}

CAPTION OR NOTES (only if the post has a caption)
{{ caption text }}

CLAIMS FOUND BY THE CONTEXT AGENT
{{ the Context agent's claims, numbered: 1. "returned 18% last year" }}

FCA RULES (ID | rule)
{{ your fca rules, one row per line }}

OUTPUT
Reply with only JSON in this shape:
{"relevant_rules":["COBS 4.2.1R","COBS 4.6.2R"],"findings":[{"quote":"exact words","rule_id":"COBS 4.6.2R","issue":"one sentence","severity":"high|medium|low","line":"L6"}]}
Copy each quote exactly as it appears in the promotion, word for word. Keep quotes short: the claim itself, not the whole sentence.
For every quote, give "line": the ID of the text line it comes from, for example "L5". If a quote runs over two lines, give the first.
In "relevant_rules", list the ID of every rule in the list that applies to this promotion, breached or not.

Agent 3: Internal rules review

Reads: Promotion, Context claims, internal rules

Same mechanics as the regulatory review, against the firm's own rulebook. This is where house style and past lessons live.

You are helping a UK financial promotions compliance reviewer. Work only from the material below. Do not invent rules, claims, IDs or facts.

TASK
Check the promotion against our internal rules below. Flag only clear breaches, one finding per issue, with the single best rule ID. Skip anything a rule does not clearly cover, and do not repeat a finding under a second rule. At most 5 findings, most serious first.

PROMOTION (text read from the image, line by line, top to bottom. size = text height as a % of the image height, so bigger means more prominent. top = how far down the image the line starts. The reader can garble very large or stylised text, so read odd characters with care.)
{{ one line per text line of the ad: L1 | text | size 6.7% | top 21% }}

CAPTION OR NOTES (only if the post has a caption)
{{ caption text }}

CLAIMS FOUND BY THE CONTEXT AGENT
{{ the Context agent's claims, numbered: 1. "returned 18% last year" }}

INTERNAL RULES (ID | rule)
{{ your internal rules, one row per line }}

OUTPUT
Reply with only JSON in this shape:
{"relevant_rules":["IR-03","IR-07"],"findings":[{"quote":"exact words","rule_id":"IR-03","issue":"one sentence","severity":"high|medium|low","line":"L6"}]}
Copy each quote exactly as it appears in the promotion, word for word. Keep quotes short: the claim itself, not the whole sentence.
For every quote, give "line": the ID of the text line it comes from, for example "L5". If a quote runs over two lines, give the first.
In "relevant_rules", list the ID of every rule in the list that applies to this promotion, breached or not.

Agent 4: Claims library check

Reads: Promotion, Context claims, claims library

Checks each claim against the claims library: approved, expired, rejected before, or not in the library at all.

You are helping a UK financial promotions compliance reviewer. Work only from the material below. Do not invent rules, claims, IDs or facts.

TASK
Match each claim in the promotion to the claims library. Check only claims that need evidence: figures and performance, awards or ratings, comparisons, sustainability, and cost. Ignore slogans and general marketing language. One entry per claim; do not split a sentence. Return only the claims that need action: not in the library, expired, or rejected before. Quote the exact words and give the library ID.

PROMOTION (text read from the image, line by line, top to bottom. size = text height as a % of the image height, so bigger means more prominent. top = how far down the image the line starts. The reader can garble very large or stylised text, so read odd characters with care.)
{{ one line per text line of the ad: L1 | text | size 6.7% | top 21% }}

CAPTION OR NOTES (only if the post has a caption)
{{ caption text }}

CLAIMS FOUND BY THE CONTEXT AGENT
{{ the Context agent's claims, numbered: 1. "returned 18% last year" }}

CLAIMS LIBRARY (today is <today's date>; a claim past its expiry date is expired)
{{ your claims library, one row per line }}

OUTPUT
Reply with only JSON in this shape:
{"claims":[{"quote":"exact words","library_id":"CL-21 or null","status":"approved|expired|rejected_before|not_in_library","issue":"one sentence","line":"L9"}]}
Copy each quote exactly as it appears in the promotion, word for word. Keep quotes short: the claim itself, not the whole sentence.
For every quote, give "line": the ID of the text line it comes from, for example "L5". If a quote runs over two lines, give the first.

Agent 5: Suggested fixes

Reads: All findings, approved wording and past decisions

Grounds each fix in approved wording, approved claims and past decisions. It drafts new wording only when nothing approved fits, and says so.

You are helping a UK financial promotions compliance reviewer. Work only from the material below. Do not invent rules, claims, IDs or facts.

TASK
Suggest one fix for each finding. Use approved wording or an approved claim first and give its ID. Follow past decisions. Draft new wording only when nothing approved fits, and say so.

FINDINGS
{{ one line per finding: R1 | "quoted words" | rule ID | issue }}

APPROVED WORDING AND PAST DECISIONS
{{ your approved wording, one row per line }}

OUTPUT
Reply with only JSON in this shape:
{"fixes":[{"finding":"R1","replace_with":"new wording, or empty to remove","source_id":"W-02 or CL-34 or drafted","note":"one sentence"}]}

Agent 6: Review report

Reads: Context output, all findings, fixes, reviewer decisions

Runs after human review. Writes up the findings, fixes and the reviewer's decisions in the standard review report format. It reports what was decided and does not review the promotion again.

You are helping a UK financial promotions compliance reviewer. Work only from the material below. Do not invent rules, claims, IDs or facts.

TASK
Write the review report. Take product, audience and channel from the Context agent. List every finding in the order given, with its fix and the reviewer's decision. Copy quotes, rule IDs, issues, fixes and decisions exactly as given. Do not reword, merge, drop or add findings, and do not review the promotion again. If a finding has no decision, write "pending".

REVIEW DETAILS
{{ date of the review }}
{{ file name of the promotion image }}

CONTEXT AGENT OUTPUT
{{ the Context agent's product, audience and channel }}

FINDINGS (ID | rule ID, library ID or not in library | quoted words | issue)
{{ every finding from the three checks, R for regulatory, I for internal rules, C for claims library: R1 | COBS 4.6.2R | "returned 18% last year." | issue }}

FIXES (finding ID | replace with, or empty to remove | source ID)
{{ the Suggested fixes agent's output, one fix per line }}

DECISIONS (finding ID | decision)
{{ the reviewer's decision for each finding: fix accepted or finding rejected }}

OUTPUT
Reply with only the report, as plain text in this shape:
FINANCIAL PROMOTION REVIEW
Date: <date>
Image: <file name>
Product: <product>
Audience: <audience>
Channel: <channel>

FINDINGS AND DECISIONS
<ID> [<rule ID, library ID or not in library>] "<quoted words>"
Issue: <issue>
Fix: remove (<source ID>)
  or: Fix: replace with "<new wording>" (<source ID>)
Decision: <decision>

One block per finding, with a blank line between blocks.

End with one line counting each decision, for example:
Decisions: 8 fixes accepted, 6 findings rejected.