

Building AI agents for multi-jurisdiction compliance

A practical workshop on agentic workflows for regulated teams

01 Introduction

02 Building an agentic workflow

03 AI impact analysis in action

04 Lessons from the field

About us



Ashi Bajwa

Regulatory AI Architect



Sam Green

Senior Policy & Partnerships
Lead



Zakaria Ouraich

Head of UK Compliance



Tim Tyler

Vice President



The anatomy of a good AI use case



High-volume / repetitive



Clear source of truth



Defined output structure



Auditable

The problem with horizon scanning



Fragmented inputs



Noise vs. signal ratio



Interpretation bottleneck



Implementation Insights

Horizon scanning as an AI use case



High volume and repetitive



Clear source of truth



Defined output structure




Auditable

POLL

POLL

**How much are you
leveraging AI for Horizon
Scanning today?**

Let's build a Horizon scanning impact agent



Case Study - Meridian Asset Management




Central monitoring across jurisdictions

Asset manager headquartered in the UK, operating across multiple countries, running central horizon scanning and routing obligations to the right jurisdictions.



A noisy feed that isn't useful

The current horizon scanning feed is high-volume and noisy. It flags a lot but surfaces little that teams can act on.



Uplift into full impact analysis

Turn that raw feed into full impact analysis, mapping each change to the impacted entities, departments and policies.

Questions




Limitations to be aware of



Configurability

Hard to scale when you are working with hundreds of updates across multiple entities and jurisdictions, and languages



Context Management

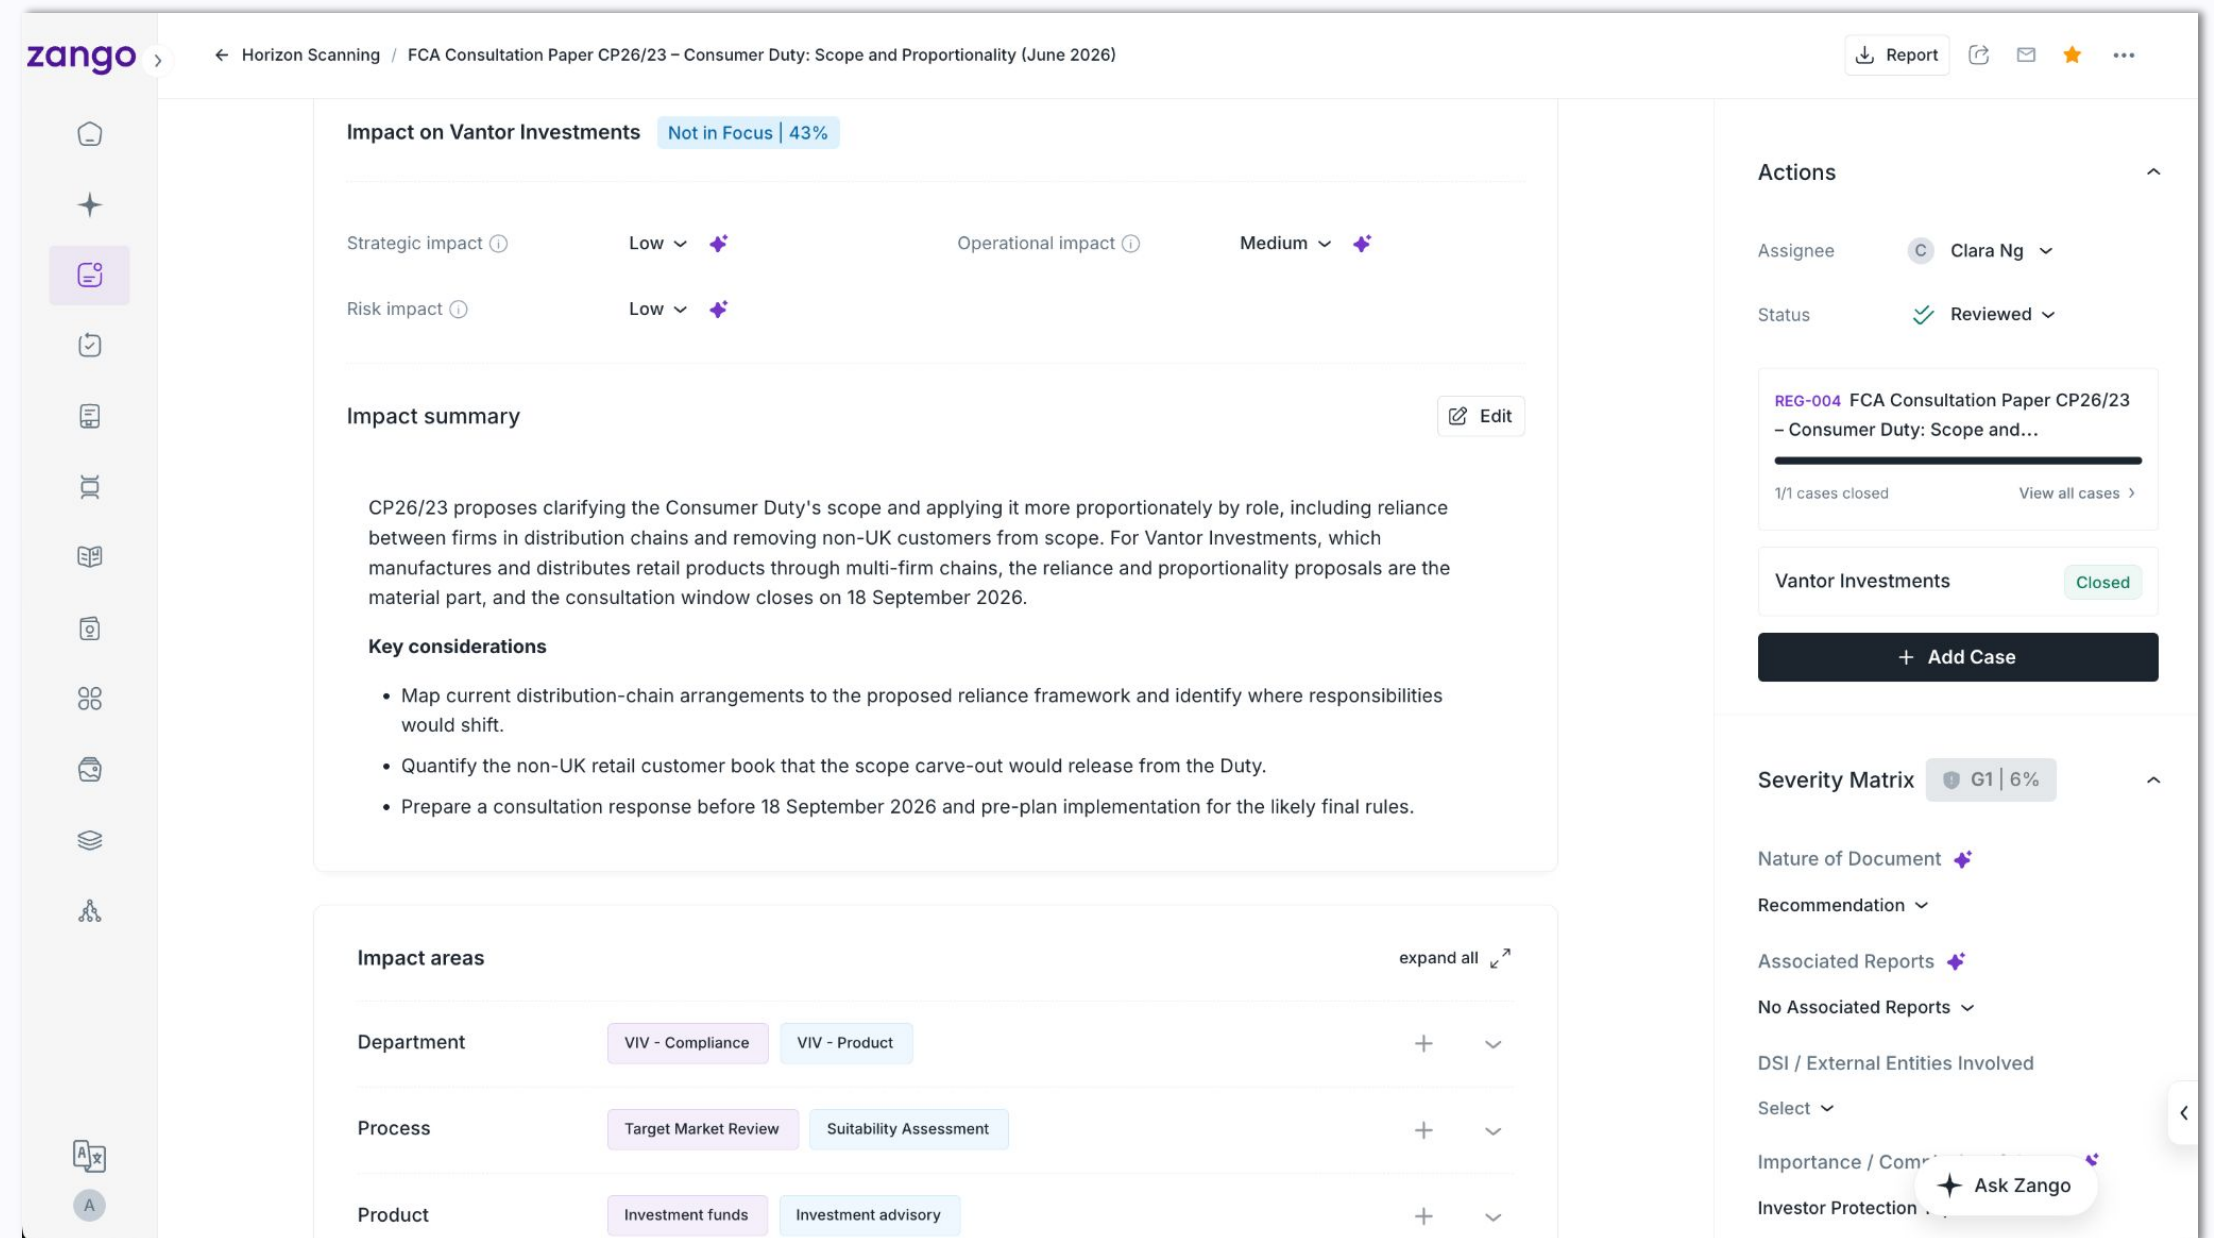
Keeping up with the organisations changing context remains very manual, as well as there being no way to embed feedback loops and reinforcement learning



Transparency and Reporting

Difficult to maintain audit records, and reporting of impact analysis, and management oversight

AI impact analysis in action



POLL

POLL

**Which agent would save
you the most amount of
time in your day to day?**

INDUSTRY INSIGHTS

Lessons from the field

Lessons from the field

01 Move beyond the chatbox

The value is in agents that run a process end to end, not one-off prompts.

02 Compliance has to get operational

The second line can no longer sit apart from the business it oversees.

03 Treat adoption as a portfolio of experiments

Run adoption as measured experiments, not a token-cost procurement exercise.

04 Compliance by design

Governance written into the build from day one, not reviewed in afterwards.

Compliance by design - KYC / AML Agent



Humans own every decision that matters

Maker-checker sign-off on every case; every screening hit goes to the MLRO.



Audit trail by default, not retrofit

Hash-chained log of every action, each input traced to source, page and quote.



Deterministic rules first, AI judgement last

Hard rules in code; AI judgement only for the ambiguous cases.



Prove it in shadow before production

Shadow-run against real cases; go-live only on documented sign-off.

Private markets compliance platform

The concept it's growing into: one platform for lean private-markets compliance teams, with six modules on a shared evidence layer, resolved once and reused across monitoring, reporting and DDQs.

KEY RECOMMENDATION

Start small, in your own function - automate the workflow you already own and can supervise.

SOME SIMPLE IDEAS

Extract mandate and fund restrictions for live portfolio compliance checks.

AI powered document classification and storage, enabling easy retrieval of documents

POLL

POLL

**What compliance AI use
case should we build
next?**

Questions



How to build AI agents for regulatory change management

Designing AI agent workflows with humans in the loop

13 Oct · 3:00 PM BST · Registrations open



Tim Tyler
Vice President



Sam Green
Policy & Partnerships Lead
zango



Ashi Bajwa
Regulatory AI Architect
zango

Contact us



Ashi Bajwa

Regulatory AI Architect, Zango

ashi@zango.ai



Sam Green

Senior Policy & Partnerships Lead, Zango

sam@zango.ai

Session plan

Slide	Slide Description	Lead	Approx Time
1	Title Slide	Tim	-
2	Agenda	Tim	-
3	About us	Tim	5 Mins
4	The anatomy of a good AI use case	Sam	2 Mins
5	The problem with horizon scanning	Ashi	-
6	Horizon scanning as an AI use case	Ashi	2 Mins
7	Poll 1	Tim	2 Mins
8	Agent Build - Title	Sam	1 Min
9	Agent Build - Objective > Transition to screenshare	Ashi	15 Mins
10	Questions	Tim	5 Mins
11	Limitations	Sam	5 Mins
12	AI impact analysis in action	Ashi	5 Mins
13	Poll 2	Tim	2 Mins
14	Lessons from the field - Title	Sam/Zakaria	-
15	Lessons from the field	Sam/Zakaria	10 Mins
16	Questions	Tim	5 Mins
17	Poll 3	Tim	2 Mins
18	Next Workshop	Sam	-
19	Contact Us	Sam	2 Mins