



## CASE STUDIES

# Governance-first AI, in the real world.

Two tracks of proof: systems we shipped and run, and Discovery engagements where AI earned its place. Working software designed to run inside your environment and owned by you, and fixed fee assessments that end in a prioritised roadmap.

Some details are abstracted to respect client confidentiality.

### BUILDS · SYSTEMS WE SHIPPED AND RUN

Working software, designed and built to run inside your environment and owned by you. One problem at a time, governance and human oversight built in from day one.

01	Signal Radar	Applied AI
02	Trend Radar	Beverage & FMCG
03	Trend Radar Global	Beverage & FMCG
04	M&A Radar	Beverage & FMCG
05	CORA	Beverage & FMCG
06	FLOID	Beverage & FMCG

### DISCOVERY · WHERE AI EARNS ITS PLACE

Fixed fee, time bound assessments. We map the operation, score use cases on value, feasibility and risk, and hand over a prioritised roadmap. No build commitment.

07	Royal Swinkels	Beverage & FMCG
08	Oldendorff Carriers	Global Shipping
09	McAlpine	Manufacturing
10	Facilities Management Group, Central Europe	Facilities & Property Management
11	RoadMender	Field Services

# Signal Radar, weekly target intelligence for commercial teams

A configurable strategic intelligence engine that turns weak signals into decision-ready opportunities, risks and assumption challenges.

Built to replace manual market scanning with a governed weekly shortlist of organisations, events and signals worth acting on.

Scope: Proven radar capability, now configured for client environments

Engagement: Designed, built, run in production

**40+**

public sources monitored

**8**

weighted signal types

**<£4**

compute per week

**5-10**

ranked targets weekly

## THE CHALLENGE

Useful market intelligence is buried across dozens of fragmented sources, and most monitoring tools drown teams in hype rather than surfacing the few signals that actually matter for a decision.

## WHAT WE DID

We built an automated radar that ingests dozens of public sources, including news, hiring, company filings, events and regulatory feeds, then applies a weighted signal framework to score organisations on genuine intent rather than noise.

## WHAT WE BUILT

A two stage language model pipeline does the work. A low cost pass extracts and tags every item, then a higher quality pass writes concise, decision ready briefs. Entity resolution merges duplicates and public company data confirms size and fit. It runs weekly on public or permissioned data only, with scoring that stays explainable.

## OUTCOME

A repeatable radar capability that converts fragmented external signals into ranked, evidence-backed decision prompts. AI Tappers now configures the same proven pattern for clients who need market, competitor, risk or opportunity monitoring aligned to their own strategic priorities.

Workflow automation

Multi-source ingestion

Two-stage LLM pipeline

Signal scoring

Company enrichment

GDPR-aware by design

# Trend Radar, MEA beverage market intelligence

A bi-weekly MEA beverage radar that turns regional market signals into a sourced intelligence brief.

Scope: Beverage market and category radar

Engagement: Designed, built, run in production

49

regional and category feeds

8

MEA-specific signal categories

24

report recipients

Bi-weekly

autonomous cadence

## THE CHALLENGE

Regional beverage teams need to track category shifts, competitor activity, launches, regulation and route-to-market movement across fragmented MEA sources. Global reporting often misses or flattens these signals, while manual scanning is inconsistent and time-consuming.

## WHAT WE DID

We configured the TrendRadar pattern for MEA beverage intelligence, combining regional news, beverage trade sources, marketing publications and targeted country, category and competitor queries. The system filters, deduplicates and classifies recent articles against MEA-specific signal categories before generating a sourced executive brief.

## WHAT WE BUILT

An n8n-based automation runs every two weeks, reading 49 regional, category and competitor feeds. A two-pass AI workflow tags and classifies articles, then generates a structured HTML report with an executive summary, trend deep-dives, competitor watch and country spotlight, delivered by email to the regional team.

## OUTCOME

A bi-weekly regional intelligence brief that keeps 24 users current without manual scanning. The same radar pattern can now be configured for other regions, categories or competitor sets by changing the source mix, signal taxonomy and reporting logic.

Beverage market monitoring

Weighted RSS sources

AI theming and summary

Regional bias control

Linked-source digest

Reusable radar engine

# Trend Radar Global, monthly beverage intelligence

A monthly global beverage radar that turns public market signals into a sourced executive trend brief.

Scope: Client build, global beverage market monitor

Engagement: Designed and built for a client, run on a schedule

## 45

global beverage feeds

## 9

predefined trend categories

## Monthly

autonomous cadence

## Linked

sourced evidence

### THE CHALLENGE

Global insights teams need a reliable read on category, competitor and market movement across fragmented trade sources. Manual scanning is time-consuming, inconsistent and difficult to scale, while generic reports often arrive too late or lack source-level evidence.

### WHAT WE DID

We configured the TrendRadar pattern for global beverage intelligence, combining industry publications and targeted news queries into a monthly monitoring workflow. The system filters, deduplicates and classifies recent articles against predefined beverage trend categories before generating a sourced executive brief.

### WHAT WE BUILT

An n8n-based automation runs monthly, reading 45 global beverage feeds and targeted news queries. A two-pass AI workflow tags and classifies articles, then generates a structured HTML report with an executive summary, radar overview, trend deep-dives, competitive moves and linked source evidence. Failure alerts notify AI Tappers if a run needs investigation.

### OUTCOME

A monthly global intelligence brief that gives the insights team a consistent, sourced view of category and competitor movement without manual scanning. The same radar pattern can be configured for different markets, categories or strategic priorities while keeping the evidence trail visible.

Global market monitoring

Weighted trade sources

AI theming and summary

Linked-source digest

Scheduled, with error alerts

Client-owned radar

# M&A Radar, acquisition signal intelligence

A weekly radar that identifies emerging beverage brands showing early signs of strategic acquisition potential.

Scope: Beverage M&A and strategic target monitoring

Engagement: Designed, built, ready for scheduled operation

## 57

market and M&A feeds

## 125

point signal framework

## 40+

alert threshold

## Verified

secondary fact-check

### THE CHALLENGE

Corporate development and strategy teams need to spot acquisition-relevant brands before they become obvious. The signals are fragmented across funding news, distribution moves, facility changes, product launches, leadership changes and category press. Manual monitoring is inconsistent, and generic market reports rarely connect these signals into a target-level view.

### WHAT WE DID

We configured a dedicated radar for beverage M&A intelligence, combining European business sources, global beverage publications, M&A feeds and targeted news queries. Instead of summarising trends, the system scores individual brands against acquisition-relevant signals and filters out large corporations to focus attention on emerging or mid-size targets.

### WHAT WE BUILT

An n8n-based workflow reads 57 public feeds, deduplicates recent articles and uses a multi-stage AI pipeline to score brands against a 125-point acquisition signal framework. Brands above the alert threshold are passed through a secondary verification step before a structured brand alert report is generated and delivered by email.

### OUTCOME

A repeatable acquisition intelligence radar that turns fragmented public signals into verified brand-level alerts for strategy, M&A and corporate development teams. The same pattern can be configured for different sectors, geographies or acquisition criteria without rebuilding the core workflow.

M&A intelligence

Brand-level scoring

Public-source monitoring

Secondary verification

Strategic target alerts

Portable n8n workflow

# CORA, Communications Optimization & Response Assistant

A governed communications assistant that drafts, rewrites, translates and QA-checks content using approved internal knowledge.

Scope: Corporate communications, multinational beverage manufacturing

Engagement: Built inside the client's Microsoft 365 environment, grounded in approved SharePoint knowledge

## EN + NL

bilingual communications

## KB-first

SharePoint-grounded outputs

## 6

pre-built communication prompts

## Human review

no automated publishing

### THE CHALLENGE

Corporate communications teams needed to produce frequent internal and external content in English and Dutch without losing brand consistency, factual accuracy or review discipline. Drafting, rewriting and translating repeat communications manually created delay, while generic AI tools risked unsupported claims, inconsistent tone and weak governance.

### WHAT WE DID

We configured CORA inside the client's Microsoft 365 environment, grounding it in an approved SharePoint knowledge base and defining clear communication tasks, tone guidance and review boundaries. CORA supports drafting, rewriting, translation and QA across common communications formats, while keeping final approval with the human owner.

### WHAT WE BUILT

A Microsoft 365 Copilot Agent Builder assistant that uses specified SharePoint sources before generating content. CORA supports English and Dutch communications, with pre-built prompts for press releases, LinkedIn posts, internal memos, complaint replies, sponsorship declines and brand-voice rewrites. The configuration is portable plain text, with no external APIs, custom code or third-party integrations.

### OUTCOME

A client-owned communications assistant that acts as a governed first-draft and QA layer for corporate communications teams, helping them produce faster, more consistent English and Dutch content while staying grounded in approved internal knowledge. Final approval remains with the human owner.

Corporate communications

SharePoint-grounded

Bilingual EN/NL

Brand-voice drafting

Human review

Microsoft 365 native

Portable configuration

# FLOID, sceptical consumer challenge assistant

A behaviourally grounded consumer persona that challenges product ideas, campaign hooks and messaging before budget is committed.

Scope: Marketing and innovation, multinational beverage manufacturing

Engagement: Built inside the client's Microsoft 365 environment using the brand's own persona and behavioural source material

## 3

explore, challenge, refine

## KB-grounded Minutes

behavioural science and brand context  
early idea pressure-test

## Human

decision, no auto approval

### THE CHALLENGE

Marketing and innovation teams often develop ideas in rooms full of category knowledge, brand optimism and internal assumptions. The harder question is whether a time-poor, low-loyalty consumer would notice, understand or act on the idea before production or media spend is committed.

### WHAT WE DID

We configured FLOID around Swinkels' own sceptical consumer persona, grounding it in behavioural science, category thinking and brand context. The assistant reviews slogans, campaign ideas, product concepts and visual prompts, then returns structured feedback designed to explore the idea, challenge the assumption and push a practical refinement.

### WHAT WE BUILT

A Microsoft 365 Copilot Agent Builder assistant inside the client's tenant, using Swinkels' persona material, behavioural science documents and decision-tree logic as its knowledge base. FLOID gives marketing teams a fast first challenge on slogans, campaigns and product ideas, with responses structured around explore, challenge and refine prompts.

### OUTCOME

A client-owned campaign challenge assistant that gives marketing teams a fast, behaviourally grounded first review before production, research or media spend. FLOID helps surface weak hooks, unclear behaviour change and lazy messaging early, while final decisions remain with the human team. FLOID is not a replacement for consumer research; it is a fast internal challenge layer that helps teams improve ideas before formal testing or spend.

Campaign pressure-testing

Behavioural-science grounded

Client persona logic

Explore / challenge / refine

Microsoft 365 native

Human decision-making

Pre-research challenge layer

# Royal Swinkels

From Discovery to a governed AI portfolio across the business.

AI Tappers helped shape a controlled portfolio of AI systems across Marketing, Corporate Communications and Global Insights: built as lean MVPs, used in real workflows, and designed for portability, ownership and future industrialisation.

Scope: Multinational beverage manufacturer, Marketing, Corporate Communications and Global Insights

Engagement: Discovery, governed architecture and multi-function AI rollout

## 4

functions supported

## 3+

AI systems deployed

## Controlled

tenant or public-source data

## Portable

no proprietary lock-in

### THE CHALLENGE

Royal Swinkels had real momentum with AI across the business, but adoption was beginning to fragment across teams, tools and licensing models. Leadership needed a governed way to scale what was working, protect data and ownership, and avoid turning early enthusiasm into long-term vendor dependency.

### WHAT WE DID

Discovery mapped where AI could create value across functions and where governance, ownership and data boundaries mattered most. We then shaped a deployment model that separated contained MVPs from future industrialisation, allowing teams to prove value quickly while preserving portability and client ownership.

### WHAT IT PRODUCED

The work produced a governed portfolio of AI systems: CORA for corporate communications, FLOID for marketing challenge and refinement, TrendRadar for global and regional market intelligence, and M&A Radar for acquisition signal monitoring. Microsoft-native assistants operate inside the Swinkels tenant, while radar systems use public-source automation with portable workflows and no hidden data persistence.

### OUTCOME

AI moved from scattered experimentation to a governed operating portfolio across Marketing, Corporate Communications, Global Insights and M&A. Teams kept the tools and workflows that worked, while Swinkels retained ownership of its prompts, sources, scoring logic and future deployment choices.

Discovery to deployment

Multi-function AI portfolio

Governance by design

Client-owned logic

Portable architecture

Lean MVP to industrialisation

“The Discovery gave us a clear view of where AI could support our teams across multiple functions. It quickly translated into deployed AI systems now supporting Marketing, Corporate Communications, M&A and regional operations. This is not an isolated AI initiative. It is becoming part of how our teams operate.”

**Sean Durkan**

Head of AI & Global Insights, Royal Swinkels

# Oldendorff Carriers

Discovery for safety-critical operations, where AI must be useful, governed and trusted before it scales.

Scope: Enterprise dry bulk shipping, operational and technical functions

Engagement: Discovery, use-case prioritisation and governed Knowledge Assistant roadmap

16

stakeholder interviews

9

structured inputs

4

function groups mapped

1

board-ready AI roadmap

## THE CHALLENGE

A global dry bulk operator had complex, distributed operations and growing interest in AI, but no shared evidence base for where to start. Leadership needed to identify where AI could reduce manual information work without disrupting safety-critical workflows, weakening accountability or introducing new operational risk.

## WHAT WE DID

We ran a structured Discovery across operations, commercial and technical stakeholders, combining interviews with role-tailored questionnaires. We mapped current workflows, surfaced pain points, and scored candidate use cases against value, feasibility, operational risk and governance readiness.

## WHAT IT PRODUCED

Discovery produced a prioritised AI roadmap and technical report leadership could act on. The recommended first build was a governed Knowledge Assistant over the company's own operational knowledge: designed for Azure, bounded by human oversight, and structured around explicit ownership before wider deployment.

## OUTCOME

A CTO-backed roadmap that turned AI interest into a controlled first deployment, with governance, ownership and operational risk addressed before build.

Safety-critical Discovery

Operational workflow mapping

Use-case prioritisation

Governed Knowledge Assistant

Azure-ready architecture

Human oversight

“The Discovery provided a clear and structured view of where AI can support our operational workflows. It reflected the complexity of our operations and identified practical opportunities to reduce manual information work and improve decision-making across teams. It provides a strong foundation for implementing AI in a controlled and operationally relevant way.”

Sönke Hoerlyk

CTO, Oldendorff Carriers

# McAlpine

Turning scattered AI experiments into a governed roadmap for adoption, risk control and internal capability.

Scope: Family-owned manufacturer, four UK sites, sales and operations

Engagement: Four-week Discovery, shadow-AI review and phased adoption roadmap

## 4 wk

structured Discovery

## 14

one-to-one interviews

## 4

sites in scope

## 2-3

quick-win pilots identified

### THE CHALLENGE

A long-established, family-run manufacturer had growing AI activity across the business, but little coordination. Teams were trialling assistants site by site, sensitive documents were entering consumer AI tools, and leadership lacked a clear view of value, risk or return. With four largely autonomous sites, uneven processes and operational data spread across ageing systems, the business needed an evidence-led view of where AI could realistically help before committing further investment.

### WHAT WE DID

We ran a four-week Discovery across sales, factory general managers and IT, combining one-to-one interviews, role-tailored questionnaires and a baseline team survey. Rather than asking teams to invent AI use cases, we mapped real workflows and sub-processes, then assessed them through a governance lens covering shadow AI exposure, data risk, GDPR and EU AI Act alignment.

### WHAT IT PRODUCED

Discovery produced an AI opportunities playbook: a prioritised view of where AI could genuinely move the needle, two to three quick-win pilots, and a phased roadmap that starts with one site before wider rollout. The recommendation was to build internal capability, governance and adoption discipline first, rather than buy more tools.

### OUTCOME

A family business with strong instincts but scattered effort got a single governed starting point: an honest sequence for building confidence, capability and risk control before larger AI investment.

Shadow AI review

Workflow mapping

Use-case prioritisation

GDPR and EU AI Act lens

Adoption roadmap

Internal capability plan

“The Discovery process helped us step back and understand where AI can realistically add value in our business. Our focus now is on building internal capability and ensuring our teams use AI effectively before investing in more advanced solutions.”

**Ross McAlpine**

Managing Director, McAlpine

# Facilities Management Group, Central Europe

Enterprise-grade AI Discovery for a complex SME operator, delivered in the client's own language.

Scope: SME facilities and property management, cost allocation, compliance, tenant reporting and field operations

Engagement: Compressed Discovery and AI Capability Catalogue, delivered in the client's language

**2 wk**

compressed Discovery

**In-language**

interviews and outputs

**9**

buildable AI modules

**4**

recommended packages

**THE CHALLENGE**

A facilities and property management group was running complex cost allocation, maintenance, claims, statutory compliance, tenant reporting and field operations through spreadsheets and disconnected tools. Leadership wanted to understand where AI could reduce manual effort without disrupting the transparent open-book model their tenant relationships depend on. Part of the team worked outside English, so Discovery had to reflect the language of the operation.

**WHAT WE DID**

We ran a compressed Discovery using role-tailored questionnaires and interviews in the client's own language. We mapped the operational workflows behind cost allocation, supplier invoices, claims, inspections, annual reconciliation, tenant reporting and field activity, then assessed where AI could reduce effort without weakening transparency, auditability or control.

**WHAT IT PRODUCED**

Discovery produced an AI Capability Catalogue: one Foundation layer and nine independently buildable modules, each separately priced, dependency-mapped and sequenced. The catalogue also gave leadership recommended packages and monthly run-rate options, so they could choose what to build, when to build it and how much to commit at each stage.

**OUTCOME**

A complex SME operator received enterprise-grade AI Discovery in its own language, translated into a practical AI build menu leadership could approve module by module. The result was a phased path from operational pain points to an owned AI operating platform, without committing to a large transformation programme upfront.

Compressed Discovery

In-language delivery

AI Capability Catalogue

Modular build roadmap

Cost allocation and compliance

Tenant open-book model

SME-fit implementation

# RoadMender

Turning operational scale-up pressure into a practical AI roadmap for growth.

Scope: UK operational scale-up, commercial, factory, field operations and sales

Engagement: AI Discovery, operational workflow mapping and 90-day implementation roadmap

**5**

stakeholder interviews

**24**

automation opportunities

**4**

operational pillars mapped

**90 days**

phased roadmap

## THE CHALLENGE

RoadMender was scaling quickly, but operational knowledge lived inside people, spreadsheets and disconnected systems. Leadership wanted to grow machine deployment and council reach without adding proportional operational load, while improving visibility across commercial activity, factory workflows, field operations and sales.

## WHAT WE DID

We ran Discovery interviews with leadership and operational stakeholders, mapping where manual reporting, individual knowledge dependency and disconnected workflows were creating scale risk. We then prioritised AI and automation opportunities across commercial operations, customer intelligence, machine monitoring, field support and leadership reporting.

## WHAT IT PRODUCED

Discovery produced a 90-day operational acceleration roadmap, sequencing quick wins, core systems and longer-term intelligence capabilities. The roadmap covered council intelligence, customer reorder alerts, CRM and dashboard foundations, delivery note and invoice automation, machine activity monitoring, fitter support and tender intelligence.

## OUTCOME

RoadMender received a practical AI implementation compass: a sequenced roadmap showing what to build first, what systems were needed underneath, and how AI could support growth without simply adding more people. The plan gave leadership and their implementation resource a clear path from manual operations to scalable operational intelligence.

Operational intelligence

AI Discovery

Workflow mapping

90-day roadmap

Council intelligence

Machine monitoring

Leadership dashboards

SME scale-up

“AI Tappers helped us step back from the day-to-day noise and see where AI could actually help RoadMender scale. The Discovery gave us a clear compass: what to build first, what to avoid, and how to turn manual processes into systems without losing the practical way the business works.”

**Harry Pearl**

CEO, RoadMender

NEXT STEP

# Start with Discovery.

Fixed fee, time bound, no build commitment. You get a clear recommendation either way. If nothing is worth building, we tell you.

[aitappers.com](https://aitappers.com) · [Start a conversation](#)