


# AI Quarterly Report

Where does the machine stop and a person take over?  
Four practitioners explain which AI use cases can run  
SMEs at scale.



# Contents

|           |                                                 |    |
|-----------|-------------------------------------------------|----|
|           | Foreword .....                                  | 03 |
|           | RUSS SHAW CBE                                   |    |
|           | About this report .....                         | 04 |
|           | The three behaviors .....                       | 05 |
| CHAPTER 1 | The expectation gap .....                       | 06 |
| CHAPTER 2 | Evidence from the field .....                   | 09 |
| CHAPTER 3 | From 150 minutes to 6 .....                     | 10 |
|           | ZOLTAN VASS                                     |    |
| CHAPTER 4 | In-moment marketing, engineered at scale .....  | 12 |
|           | GANAPATHY SANKARABAAHAM                         |    |
| CHAPTER 5 | Two agents before five .....                    | 14 |
|           | TSVETELIN ANASTASOV                             |    |
| CHAPTER 6 | From CRM entropy to customer intelligence ..... | 16 |
|           | ALEX FIORE                                      |    |
| CHAPTER 7 | Key takeaways .....                             | 18 |
|           | About the GTA AI Quarterly .....                | 21 |

IN THIS REPORT

Four practitioners from the Global Tech Advocates network describe AI systems they have already put into production for small and mid-sized companies. Every number was reported by the person who built the system. The report closes with what they agree you should do next

## ■ EXECUTIVE SUMMARY

# Foreword

**Russ Shaw CBE**

FOUNDER, TECH LONDON ADVOCATES &  
GLOBAL TECH ADVOCATES



**When we founded Tech London Advocates, and subsequently Global Tech Advocates, the mission was clear: to build an inclusive, private-sector-led network that connects tech ecosystems, champions innovation, and ensures technology acts as an engine for economic growth.**

Today, across our network, which spans more than 45 hubs across the world, our role as a convener is more critical than ever. Technology never develops in isolation, and this report represents a direct step toward fulfilling our mission, surfacing on-the-ground insights from inside our community and turning them into actionable insights.

Over the past few weeks, the global conversation around artificial intelligence has swung dramatically toward fear and doomsday headlines. While headlines suggest panic over market volatility and high-profile resignations, the view on the ground for businesses tells a very different story. What I am hearing directly from our tech leaders, founders and working groups is that AI has firmly moved past the speculation and into real-world commercial deployment.

AI is now proving to be a great economic leveller. Across a wide variety of ecosystems, advocates are using these tools to solve practical operational problems, saving time, boosting productivity, and sparking an explosion of startups who are building and scaling businesses from day one.

To capture these ground-level dynamics as they unfold, I am delighted to introduce this publication as the first in a recurring series produced in partnership between Global Tech Advocates and AI Central.

By pairing GTA's international reach with AI Central's deep sector expertise, this series will track the real-time shifts in founder sentiment, highlight emerging deployment trends, and provide business leaders and policymakers with the clear-sighted data needed to navigate this shift.

As the findings in this report demonstrate, there is every reason to remain confident in what lies ahead, and I look forward to tracking this journey together.

RUSS SHAW CBE, FOUNDER, TECH LONDON ADVOCATES & GLOBAL TECH ADVOCATES

# About this report

The GTA AI Quarterly session brought together members of the GTA from across the globe to discuss progress, explore emerging AI capabilities, and mull over strategic priorities for effective AI implementation. Designed as a collaborative forum, the session highlighted key developments in AI initiatives, showcased practical use cases, and encouraged the exchange of knowledge and best practices across GTA. It also provided an opportunity to discuss industry trends, assess new technologies, and identify opportunities to accelerate innovation while maintaining responsible AI governance. By fostering cross-functional collaboration and regular dialogue, the GTA AI Quarterly is designed to help GTA members ensure that AI investments remain aligned with business objectives and deliver measurable value across their organizations

## IN THIS EDITION

Access to AI stopped being the constraint some time ago. The gap now sits between what leaders believe they understand and what they can show for it, with McKinsey reporting that 96% of CMOs are excited about AI while 80% see risk to their own role. Our first GTA AI Quarterly brought together 3 practitioners who have built past that gap at SME scale. Zoltan Vass cut purchase order processing at a 200-person European wholesaler from 150 minutes to 11 minutes, at \$0.15 per document. Ganapathy Sankarabaaham turned a daily UK government licensing list of 140,000 companies into same-day qualified leads for a law firm, adding 5-10% revenue growth. Tsvetelin Anastasov talked a real estate broker out of a 5-AI-agent system and down to a simpler one with 2, ensuring a responsible and accountable build and lifting realtors' pipeline coverage from 38% to near-complete. Each made the call about where the machine stops and a person takes over. Collectively, they see today's AI systems as copilots that augment humans' capabilities and capacity, rather than replace them

## CONTRIBUTORS



### Zoltan Vass

AI TRANSFORMATION AND  
DELIVERY DIRECTOR,  
WWW.FUTUREOFWORK.UK  
· GTA FUTURE OF WORK,  
CO-FOUNDER

[linkedin.com/in/zvass](https://www.linkedin.com/in/zvass)



### Ganapathy Sankarabaaham

FOUNDER AND CEO, XITE  
CREATE


[linkedin.com/in/ganapathy-  
sankarabaaham](https://www.linkedin.com/in/ganapathy-sankarabaaham)



### Tsvetelin Anastasov

CO-FOUNDER, THINK  
REVOPS · GTA DATA AND  
AI CO-LEAD

[linkedin.com/in/tsvetelinanas-  
tasov](https://www.linkedin.com/in/tsvetelinanastasov)



### Alex Fiore

FOUNDER AND CHIEF  
EDITOR, AI CENTRAL  
MEDIA

[linkedin.com/in/alex-ai](https://www.linkedin.com/in/alex-ai)



### Marc Duke

MODERATOR ·  
COMMUNITY MANAGER,  
TLA CREATIVE TECH

[linkedin.com/in/marcduke](https://www.linkedin.com/in/marcduke)

# The three behaviors

What separates the top 10% of adopters from the other 90% is organizational, not technical

## 01

### Knowledge sharing across departments

In any organization of more than 10 or 15 people, some employees are more AI literate than the rest. What these champions discover applies horizontally: the person in admin can teach the person in sales, sales can teach HR. In most companies that knowledge stays siloed, and the company pays a consultant to explain what one of its own people is already doing

**"Your employees don't talk to each other, your employees don't share best practices and knowledge gets siloed across departments"**

#### ■ IN PRACTICE

**AI Ignite sessions** - one-day workshops that end with something built

**AI Thursdays** - cross-functional teams take one problem for 2 to 3 weeks. Non-coders built an image optimization tool now in production

**Keep it employee-driven** - colleagues inspiring colleagues is the fastest path

## 02

### A 1% daily gain, measured against time

Instead of aiming at 30% more productivity, gain 1% and make it stick every day. Executives read the balance sheet; employees live in their daily workflows, and the two rarely talk. If people are not using AI every day, the headline benefit never reaches the balance sheet. Productivity is output over time, and most companies only track one of the two

**"AI adoption is not done by the CEO most of the time. It's done by your employees"**

#### ■ IN PRACTICE

**A check-in form 3 times a day** - morning, afternoon and close, feeding a Notion database, so output sits next to time

**Self-reporting is the zero-cost option** - not a vibe check, and not a count of emails sent

**Show the progression** - once people see their own momentum, nobody wants to stop

## 03

### Executives who own build versus buy

Corporate finance 101 rather than anything specific to AI, and still the decision that separates the two groups. Some software can now be replaced, some licenses cut or rearranged. Give your most AI-literate people a little coding and a small internal hackathon, and they may build the thing you were about to buy. The best talent leaves organizations that will not let it build

**"Some of the times, the best talents in an organization, they go away because their leaders, they don't allow them to build"**

#### ■ IN PRACTICE


**Review every SaaS license** - against what an internal build now costs

**Run small internal hackathons** - and let people MVP what they need

**Treat building as retention** - talent is the scarcest commodity in the market

#### ■ THE NON-DISRUPTIVE WAY IN

The non-disruptive way in is a sandbox: keep the safe workflow running while people experiment next to it. Once one colleague closes quota earlier with AI, nobody wants to be the person who doesn't. Copycats are the fastest adoption program there is, and it costs nothing



# The expectation gap

# The expectation gap

# 1

There is a strange dichotomy at the top of most companies right now

Leaders are excited, and leaders are anxious, often the same leaders at the same time. McKinsey's latest marketer survey, run with Google, puts hard numbers on it - almost every CMO is excited, most are anxious, and four in five see a risk to their own role. Nearly 60% of marketers already use AI several times a week. Fewer than one in ten organizations capture value across a full workflow

Excitement is high. Deployment is high. Value is low. That is the gap

Practitioners are aware of what AI can and cannot do. The frictions live often with the decision-makers who must adopt AI solutions at scale

## ■ EXHIBIT 1

### CMO sentiment versus captured value

**96%**

report high excitement with AI

**71%**

express anxiety about it

**80%**

perceive a risk to their own role

**<10%**

of organizations capture value across a full workflow

Source: McKinsey and Google, Global Marketer Survey, published July 2026, about 500 marketers

Vendors see the expectation gap most clearly in the pricing conversation. Ganapathy has been in the room when a client reads a report, sees a headline number, and tries to renegotiate

**“You have quoted X on the project. I have read a report which said AI is able to automate 75 to 80%. Can you do this project for 15 to 20% of the quoted price?”**

That is the expectation gap in one sentence. The headline says 80%, so the buyer assumes the price should drop 80%, and the actual work of scoping, cleaning data, and keeping a human in the loop disappears from view

The market has already run this experiment at scale. In 2024 a wave of AI SDR tools promised fully automated outbound. Tsvetelin’s team tested two or three of them. “They were sold as a finished product and they were still very much in MVP stage,” he said. Within 12 months revenue leaders had gone back to the co-pilot model, with a person signing off every draft before it ships

# 2

The 2024 reset - tools sold as finished autonomous products shipped as MVPs. The market corrected back to humans signing off on AI-drafted work, and that is where the durable value is


■ EXHIBIT 2

## Four ways to bring an agent into the business

|    |                                                                                                                                                                 |              |
|----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| 01 | <b>Adopt the agentic features of software you already run</b><br>HubSpot, Salesforce and the rest of the stack have shipped their own agents. A vendor you know | LOWEST RISK  |
| 02 | <b>Adopt an AI-first tool</b><br>Built with agents at its core, from a new vendor you have never worked with                                                    | HIGHER RISK  |
| 03 | <b>Build in-house</b><br>The riskiest option, and the one that asks the most of the team you already have                                                       | HIGHEST RISK |
| 04 | <b>Go hybrid</b><br>Keep the SaaS you know for the core and pilot specific features of the AI-native tools where it falls short                                 | TWO LANES    |

Whichever path a company takes, Ganapathy pointed out, the decision is not really a technology choice. It is a question of how data flows and integrates across all of those products. Practitioners do not have this problem. When you talk to the people building the systems, they know exactly what AI can and cannot do.

Zoltan saw the fix in his own project. Explaining the agent changed no minds. “When I built the agent and I showcased it at demo, that’s when they started actually understanding,” he said. Closing the gap is not a training problem you solve with a slide deck.



# Evidence from the field

## ■ USE CASE 1

# From 150 minutes to 6

Power Belt, a European wholesaler, cut the processing of a complex supplier document from around 150 minutes to six by giving the matching step to an AI agent, with human review, rather than a rules-based bot



**Zoltan Vass**

AI TRANSFORMATION AND DELIVERY  
DIRECTOR · FUTUREOFWORK.UK  
GTA FUTURE OF WORK CO-FOUNDER AND  
CO-CHAIR  
[linkedin.com/in/zvass](https://www.linkedin.com/in/zvass)

## ■ THE SITUATION

Power Belt is a European wholesaler with around 200 employees. Between 700 and 800 supplier documents arrive every month, largely as PDFs by email. A complex document could take around 150 minutes to process manually. 90% or more of the effort went into manual matching and checking: matching supplier product codes to the internal product ID list

## ■ WHAT WE BUILT

The agent reads the PDF, extracts product codes, quantities and line details, and checks the supplier codes against the internal product ID list. It flags missing, conflicting or unclear information with a reason. Where an ID is ambiguous or duplicated, it does not guess. It flags the line and stops. A person resolves it and approves the output

## ■ HOW IT RUNS

On a 253-line supplier document, agent processing took 58 seconds. It matched 250 lines and flagged 3 with reasons for human review, with zero guesses. Total handling time, including human review, was six minutes. Agent compute costs around £90 a month for 750 supplier documents

## ■ WHAT IT CHANGED

Processing of a complex document went from around 150 minutes to six, a 96% reduction. For 750 documents a month, around £5,850 of human matching time becomes around £90 of agent compute, and no order enters the system without a person approving it. Automation should stop where confidence ends. The agent does not chase suppliers, update the reference table, or route for approval. It flags ambiguous cases for human review, and that boundary is what makes it trustworthy

**“The real challenge is no longer access to AI. It is to stop automation where confidence ends. That is how we build trust in AI while keeping humans in control”**

## ■ INSIGHT

Rule-based bots only match codes covered by predefined mappings. One can read “5 pallets” as 5 pieces, or struggle with a duplicate code that points to more than one product. The agent checks the product ID list and flags what it cannot resolve, with a reason

■ USE CASE 1 · THE EVIDENCE

# Where the time went

■ EXHIBIT 3

## Handling time for a complex supplier document, minutes

BEFORE


### Manual matching

Around 150 minutes to process a complex document by hand, with 90% or more of the effort on matching and checking

AFTER

### Agent processing

58 seconds of agent processing, then human review of the 3 flagged lines and approval, six minutes in total



■ EXHIBIT 4

## How the agent runs

|    |                 |                                                                                                                                                                                               |
|----|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 01 | <b>Extract</b>  | The agent reads the emailed PDF and pulls product codes, quantities and line details                                                                                                          |
| 02 | <b>Match</b>    | Supplier product codes are checked against the internal product ID list, the step that took 90% or more of the manual effort                                                                  |
| 03 | <b>Validate</b> | Missing, conflicting or unclear information is identified before anything reaches the wholesaler's system                                                                                     |
| 04 | <b>Flag</b>     | Clear lines are matched. Duplicate candidates, unrecognised units and other unclear lines are flagged with a reason, not guessed. On the 253-line document: 250 matched, 3 flagged, 0 guesses |
| 05 | <b>Review</b>   | A person resolves the flagged lines, reviews the completed output, approves or corrects it and enters the information into the system                                                         |

A quantity given in pallets is a good example. A rules-based bot can read "5 pallets" as 5 pieces. The agent did not recognise the unit, so it flagged the line with a reason rather than guessing, and a person resolved it

The scope boundary was deliberate. Full agentic workflow, where the system chases the supplier and updates the reference table itself, was not built. Resolving a duplicate supplier code is a supplier data problem, not an automation problem, and automating around it would have hidden the cause

On the 253-line document, 3 lines, around 1%, were flagged as exceptions. The exception rate is the practical measure to watch. It tells you whether the reference data is improving, and it is a more honest signal of progress than the headline time saving

**"We are able to work more efficiently, and new opportunities will open up for various positions. I think AI is a great technology which, if used properly, can work to our advantage."**

OLIVER KNAPP, COO, POWER BELT

## ■ USE CASE 2

# In-moment marketing, engineered at scale

**We built an automated pipeline that turned a daily government licensing list into same-day, qualified leads, so our customer never missed a revenue opportunity**



**Ganapathy Sankarabaaham**

FOUNDER AND CEO, XITE CREATE

[linkedin.com/in/ganapathy-sankarabaaham](https://www.linkedin.com/in/ganapathy-sankarabaaham)

## ■ THE SITUATION

Our client is a UK law firm, in business for over two decades, looking to increase revenue through better lead generation. The companies listed on a government licensing portal are exactly their target market. The portal publishes this list every day, but there are over 140,000 companies on it. Spotting new licence registrations, suspensions and revocations inside that long list, by hand, was not realistic

## ■ WHAT WE BUILT

We built a pipeline that pulls the list daily, compares it against the previous day's file, and buckets every company as new, suspended or revoked. Limited government data was then enriched using AI, adding email addresses, founder's phone number, heads of legal or finance, and LinkedIn details, so each entry became a company the firm could actually approach

## ■ HOW IT RUNS

Every morning, the AI engine flags the day's additions and deletions, ranks each company by opportunity, and drafts an email sequence and a call script with the research and value proposition attached. A salesperson only reviews and sends. What used to be days of manual list-checking now runs before the sales team logs on

## ■ WHAT IT CHANGED

Through 'in-moment marketing', reaching companies on the day their status changed, while the news was still fresh, increased the response rate. New registrations welcomed the outreach, revoked licences wanted help fighting to get reinstated, and the firm signed new clients it would never have found by chance, adding 5-10% revenue growth

**"Every SME is sitting on data that could be a revenue stream. AI is simply the tool that converts that data into daily insights for growth."**

GANAPATHY SANKARABAAHAM

## ■ INSIGHT

Regulatory and licensing data are public, many are updated daily, but almost never acted on in real time. For firms willing to build the pipeline once, in-moment marketing turns routine government publishing into a steady, repeatable source of qualified leads

■ USE CASE 2 · THE EVIDENCE

# From list to leads

■ EXHIBIT 5

## Before and after

BEFORE

### Manual list-checking

Someone had to scan 140,000 records by hand every day, if at all

AFTER

### Automated, same-day outreach

New, suspended and revoked companies flagged, enriched and pitched within hours of publication

140,000

companies on the register, republished every day

Same day

from a change of status to a drafted, reviewed approach

5-10%

revenue growth from clients the firm would never have found by chance

Figures as reported by the practitioner

■ EXHIBIT 6

## The daily pipeline

|    |                |                                                                                                     |
|----|----------------|-----------------------------------------------------------------------------------------------------|
| 01 | <b>Extract</b> | Custom code downloads the daily licence list and compares it to the previous day's file             |
| 02 | <b>Bucket</b>  | Compare against existing CRM information for better lead qualification                              |
| 03 | <b>Enrich</b>  | Data tools add emails, founders, finance and legal contacts, and LinkedIn profiles for each company |
| 04 | <b>Rank</b>    | Each lead is scored and qualified against the firm's ideal client and likely needs                  |
| 05 | <b>Reach</b>   | AI drafts an email sequence and a call script, ready for the sales team to send                     |
| 06 | <b>Learn</b>   | AI learns based on successful connects/patterns                                                     |

The idea began with one observation: a government list published every day, sitting there, unused. From that single insight, we stitched together custom code, enrichment tools, lead scoring and automated outreach into a closed-loop system that runs itself, turning public data into a private pipeline of qualified prospects for the firm

The value came from matching automation to the moment. A company that had its licence revoked yesterday is thinking about it today. If we reach them a week later, the conversation has already moved on. Speed and relevance, together, turned public data into revenue

## ■ USE CASE 3

# Two agents before five

A real estate broker asked for 5 AI agents; we shipped two, set the foundations, proved the model, and left the other 3 for six months later



**Tsvetelin Anastasov**

CO-FOUNDER, DATA, AI AND GROWTH ·

THINK REVOPS

GTA DATA AND AI CO-LEAD

[linkedin.com/in/tsvetelinanastasov](https://linkedin.com/in/tsvetelinanastasov)

[thinkrevops.com](https://thinkrevops.com)

## ■ THE SITUATION

The client is a real estate broker, running 20+ realtors handling roughly 650 inbound leads per month. That's around 30 new leads per month. Leads stay commercially alive for 3 to 6 months, which brings the total leads realtors have to engage with at once to between 90 and about 180

## ■ WHAT WE BUILT

They wanted five AI agents. We scaled it to two, with the goal to build foundations first, prove the models work, and build the rest after six months. A researcher agent takes an investor's name, scans criteria and sources, including net worth signals, countries of presence and property held elsewhere, then writes an enriched profile back to the CRM. A second agent drafts the outreach messages and sequences

## ■ HOW IT RUNS

Our discovery put each agent at 15hrs of monthly pipeline work against 39hr needed. The AI stack takes the 15hrs down to about 4.5hrs, which are mostly review of outputs. Every draft goes to a human for a sign-off and is routed through WhatsApp rather than a dashboard, to gain approval speed

## ■ WHAT IT CHANGED

The hour reduction matters less than the leads coverage increase. Pipeline touched on cadence goes from roughly 38% to near complete, and first contact drops from days to minutes. The lesson is to lead with the coverage gap rather than the cost savings. Agents were doing 15hrs of a 39hr job, missing on potential deals

**"Companies and revenue leaders realised that you cannot not have the human in the loop at the end signing off the draft"**

TSVETELIN ANASTASOV

## ■ INSIGHT

A RevOps survey has leaders scoring themselves 7/10 on AI knowledge, yet only 1 in 10 report actual ROI. McKinsey's marketing data explains part of that gap. 96% of CMOs report excitement, 71% anxiety, and 80% see risk to their own role

■ USE CASE 3 · THE EVIDENCE

# Closing the lead coverage gap

■ EXHIBIT 7

## Before and after

BEFORE

### Half-worked pipeline

Each realtor gives 15 hours to a job that requires 39 hours, every month

AFTER

### Full-cadence pipeline

Realtors' time drops to 4.5 hrs; coverage rises from 38% to near complete

### 600 hours

Monthly pipeline work across all real estate agents going untouched

### 240 hours

Returned to client-facing work each month, roughly 1.5 FTE real estate agents

### 30% to 100%

Share of leads arriving pre-researched and tiered by the AI agent

### Days to minutes

Time from lead landing to first human contact

Through discovery, we were able to shift the use case from pure cost-savings play to coverage increase and productivity gains

■ EXHIBIT 8

## The two-agent pipeline

|    |            |                                                                                                                                       |
|----|------------|---------------------------------------------------------------------------------------------------------------------------------------|
| 01 | Data audit | Assess the CRM and other relevant company data quality first, because the quality of the inputs impacts the quality of the output     |
| 02 | Input      | The researcher agent receives the investor's name, initial profile, and personal details from the CRM or enrichment layer             |
| 03 | Scan       | The agent checks criteria including net worth signals, countries of presence and property owned elsewhere                             |
| 04 | Profile    | Findings are assembled and enriched into a single investor profile written back to the CRM                                            |
| 05 | Draft      | The outbound agent writes the message based on research and found signals, and a human reads and signs it off before anything is sent |

There are 3 ways to get agents into a business. Use the agentic features in tools you already run, like HubSpot or Salesforce. Buy from an AI-first vendor built around agents from the start. Or build in-house, which is the riskiest of the 3 and the one that tends to proliferate and create shadow AI/IT

None of this works if the team feels unsafe experimenting. We started running lunch-and-learns at a client a few years back and now push clients toward active exploratory group sessions in the form of workshops, where a mixed group spends a day with a tool and walks out having built something practical. The goal is to inspire them to keep going and identify the gaps where AI systems can bring tangible value

On AI job displacement, I'm optimistic. Ten years ago, a churn prediction model needed machine learning specialists, dedicated budget and project management. Now an early-stage startup with a limited budget can build an early warning system on support data as a quick win. Marketing and revenue teams are progressively freed from low-level technical, analytical and reporting work to more strategic, higher-value activities and work that they just didn't have time to address before and was slipping away

## ■ USE CASE 4

# From CRM entropy to customer intelligence

A 300,000-reader media company unified data from three platforms and several domains into one customer view, then let AI optimize revenue streams



**Alex Fiore**

AI CENTRAL · FOUNDER AND CHIEF EDITOR

[linkedin.com/in/alex-ai](https://linkedin.com/in/alex-ai)



## ■ THE SITUATION

AI Central grew from zero to 300,000+ readers across beehiiv, LinkedIn and Substack, plus several domains and apps. Every platform held its own version of the reader. The CRM said “email sent” and not much else. Nobody could say who the best readers were, what each cohort was worth, or which of them would buy

## ■ WHAT WE BUILT

We unified the data before touching any AI. Every reader record was enriched and matched to behavior across sites, apps and inboxes, then bucketed into dynamic cohorts: superfans (lifetime open rate above 70%), engaged readers (click rate above 2%) and ghosts (no opens in 3 months). Lifetime value and acquisition cost sit against each cohort

## ■ HOW IT RUNS

Segments recalculate on their own as behavior changes. AI drafts the sequences, picks the offer for each cohort and writes the follow-ups. A person approves before anything ships. The team is 5 to 6 people. Productivity is measured with a check-in form 3 times a day feeding a Notion database, so output always sits next to time

## ■ WHAT IT CHANGED

The CRM turned into customer intelligence: one profile per reader across email, LinkedIn, WhatsApp, Telegram and paid ads, and a funnel that sells to 100 people who matter rather than 1,000 who don't. The lesson travels: the agent is only as good as the data underneath it, and CRM entropy is the number one bottleneck we meet with clients

**“I don’t want my salespeople to send 1,000 emails. I want to send maybe 100 good emails that sell”**

## ■ INSIGHT

McKinsey's 2026 marketer survey found nearly 60% of marketers use AI several times a week, yet fewer than 10% of organizations capture value across a full workflow. The model is rarely the problem. The data underneath it usually is

Source: McKinsey and Google Global Marketer Survey, published July 2026, about 500 marketers

## ■ USE CASE 4 · THE EVIDENCE

# One view of the customer

## ■ EXHIBIT 9

## Before and after

## BEFORE

### A badge

The CRM says “email sent” and nothing else, activity without understanding

## AFTER

### A real picture

Email, LinkedIn, WhatsApp, Telegram and paid ads connected into one profile per reader

## ■ EXHIBIT 10

## The numbers

# 300,000+

readers unified across 3 platforms and several domains

# 3

dynamic cohorts, recalculated as behavior changes


# 3x a day

check-in form feeding the productivity database

Two of the last three data projects we ran for clients were about nothing but the foundation. If the data is a mess, no enrichment tool can match it and no agent can act on it. For SMEs under \$5 million in revenue this is the first use case they fall in love with: knowing who is on your website

What separates the top 10% of adopters from the other 90% is organizational, not technical. Knowledge sharing across departments, so the AI champion in admin teaches the person in sales. A 1% daily gain that sticks, measured against real time tracking, instead of a 30% headline. And executives who are good at build versus buy

The non-disruptive way in is a sandbox: keep the safe workflow running while people experiment next to it. Once one colleague closes quota earlier with AI, nobody wants to be the person who doesn't. Copycats are the fastest adoption program there is, and it costs nothing



# Key takeaways

## ■ KEY TAKEAWAYS · 1 OF 2

**Zoltan Vass**

AI TRANSFORMATION AND  
DELIVERY DIRECTOR,  
WWW.FUTUREOFWORK.UK · GTA  
FUTURE OF WORK, CO-FOUNDER

[linkedin.com/in/zvass](https://www.linkedin.com/in/zvass)

## "Automation should stop where confidence ends"

Zoltan rebuilt the purchase order processing workflow for a European wholesaler with around 200 staff and 700 to 800 orders a month. An AI agent extracts, matches and validates each PDF, then flags anything ambiguous for human review. Processing fell from 150 minutes to 11, at \$0.15 per document. The design decision that mattered was where to stop: the agent flags exceptions and hands them to a human rather than resolving them alone and solving problems by assumptions

### WHAT TO DO NOW

- Time each step, not the whole process
- Find the step that consumes most time
- Decide where the agent flags and stops
- Track exception rate, not just time saved

**Ganapathy Sankarabaaham**

FOUNDER AND CEO, XITE  
CREATE

[linkedin.com/in/ganapathy-sankarabaaham](https://www.linkedin.com/in/ganapathy-sankarabaaham)

## "True competitive advantage now belongs to SMEs that operationalise AI faster than their peers"

Ganapathy Sankarabaaham, Founder & CEO of XITE Create, shows how a single publicly available data list became a scalable revenue engine for a UK law firm. By combining custom data extraction, enrichment, and AI-driven outreach, his team turned a daily government publication into same-day, qualified leads. The result was measurable revenue growth, and a reminder that the biggest opportunities are often hiding in data everyone can already see, waiting for someone to act on it at the right moment.

### WHAT TO DO NOW

- Identify a pain point (revenue generation / cost optimisation)
- Be creative when defining the solution
- Build a closed loop system, that learns and improves
- Measure and optimise

## ■ KEY TAKEAWAYS · 2 OF 2

**Tsvetelin Anastasov**CO-FOUNDER, THINK REVOPS ·  
GTA DATA AND AI CO-LEAD[linkedin.com/in/tsvetelinanastasov](https://linkedin.com/in/tsvetelinanastasov)

## "Let's start with 2 agents first and prove they work"

Tsvetelin is arguing that starting small and proving a concept wins in the long term, as this model helps get stakeholders on board more easily with tangible and tracked results. Allowing the organisation to mature when implementing a new technology or way of work is also critical for success. A client wanting a full-blown multi-agent architecture with 5 agents, was convinced to start with 2, on the condition that the foundations are built first and proved out before the rest was built. His broader point is that the revops AI market has come back to copilots and away from fully autonomous systems, after the early 2023-2024 AI SDR tools underperformed, and that human sign-off is now the default rather than a compromise

---

**WHAT TO DO NOW**

- Identify the real gap in the current process
- Audit internal data before building and plugging the agents
- Route every output through human sign-off
- Review, adjust, hit milestones and build further based on proven results

**Alex Fiore**FOUNDER AND CHIEF EDITOR  
AI CENTRAL MEDIA[linkedin.com/in/alex-ai](https://linkedin.com/in/alex-ai)

## "It's not fine to pretend to have clarity"

Alex Fiore unified reader data across beehiiv, LinkedIn, Substack and several domains into one customer view for AI Central's 300,000+ readers, then used AI to lift conversions on top of that base. His argument is that the constraint is organizational: fix the data foundation before the agent, share AI wins and failures across departments, chase a 1% daily gain measured against honest time tracking, and let executives own the build versus buy call

---

**WHAT TO DO NOW**

- Fix the data foundation before the agent
- Track time, then measure output against it
- Share AI wins and failures across departments

GTA AI QUARTERLY • OCTOBER 2026

# Where does the machine stop and a person take over?

## ABOUT GLOBAL TECH ADVOCATES

Global Tech Advocates is an international, volunteer-led, not-for-profit grassroots tech community that unites over 40,000 business leaders, experts, and investors across more than 60 locations worldwide.

## ABOUT AI CENTRAL MEDIA

AI Central Media is a media and e-learning company reaching 300,000+ senior professionals across North America and Europe, with a boutique AI advisory for SMEs. Its editorial team is based in London and led by Alex Fiore, chief editor and founder