DEMANDBASE

eBook

The agentic AI GTM playbook

Real-world strategies for deploying AI agents across sales, marketing, and operations

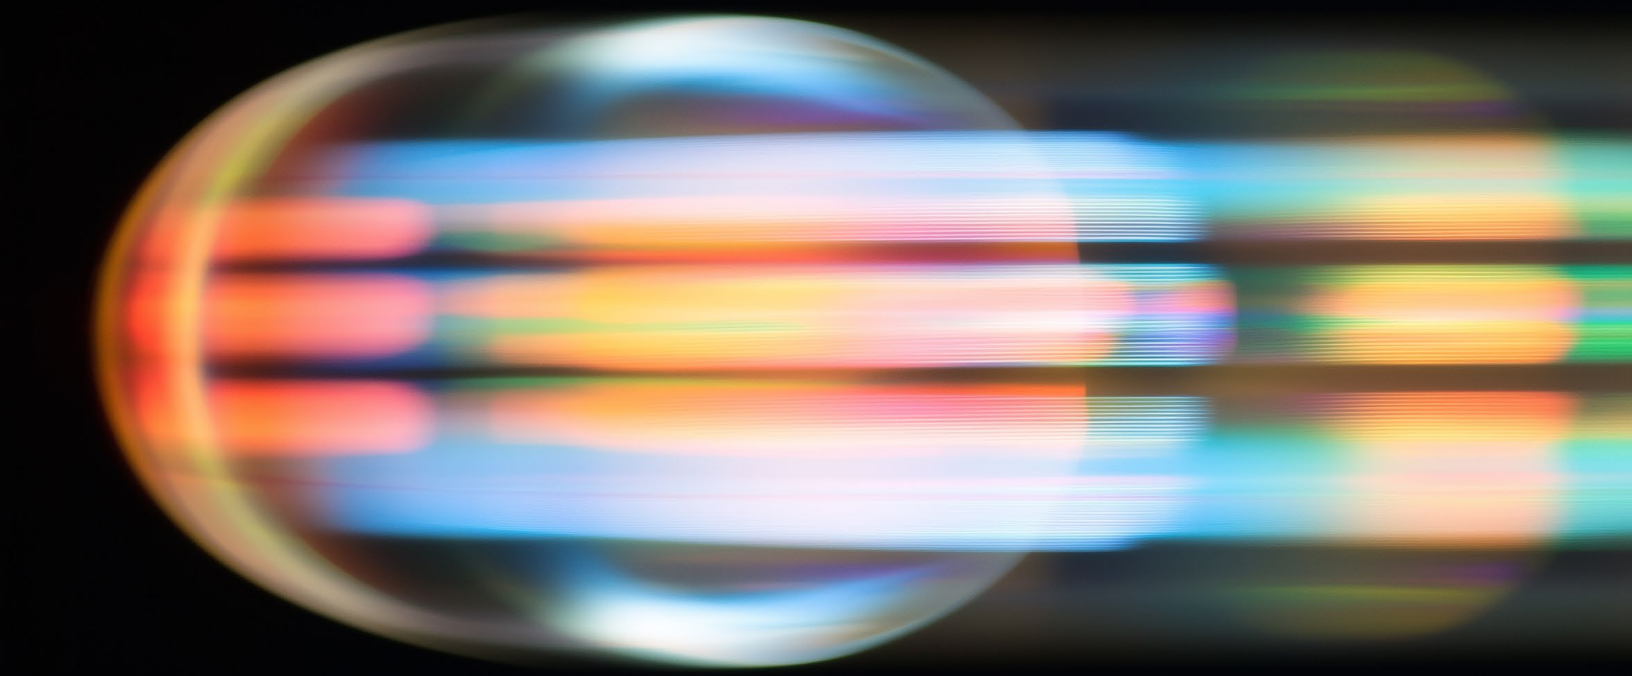

EBOOK

Table of contents

The shift to agentic go-to-market

There's a new kind of GTM motion emerging, and it doesn't just rely on data, automation, or segmentation. It runs on intelligent agents.

Agentic AI isn't just tech jargon, it represents a structural shift in how sales, marketing, and customer teams operate. It replaces static lead scoring with dynamic, always-on signal recognition. It turns one-size-fits-all campaigns into hyper-personalized, continuous conversations. And it enables teams to move faster at scale without sacrificing insight or control.

This playbook is built from insights shared by more than a dozen B2B leaders on the OnBase podcast. These aren't theoretical perspectives, they reflect how teams are already deploying AI to drive pipeline, improve execution, and scale GTM impact.

Across industries and use cases, a consistent pattern emerges: Agentic AI works when it's embedded into workflows, not layered on top of them.

You'll hear from:

- CEOs scaling outbound through signal stacking and AI clones
- ABM leaders building 1:1 campaigns without manual handoffs
- Product strategists launching real AI agents inside platforms
- Marketing ops leaders orchestrating personalization at scale

This playbook is designed to be practical. Inside, you'll find:

- A framework for deploying agentic AI across the GTM stack
- Real-world use cases already running inside B2B teams
- Quotes, strategies, and checklists from experienced operators
- Pilot-ready prompts and planning tools

Whether you're just getting started or scaling adoption across teams, the goal is the same: make Agentic AI actionable, measurable, and tied to business outcomes.

But to do that, you need more than use cases. You need a model.

The Agentic GTM operating model

Agentic AI isn't a feature. It's a new operating model for go-to-market.

Most teams begin by adding AI into existing workflows, automating tasks, generating content faster, or improving efficiency at the margins. But real impact comes from redesigning how GTM works end-to-end.

The highest-performing teams aren't just using AI. They're building systems that can continuously sense, decide, and act.

We call this the Agentic GTM Operating Model.

The five layers

1. Signals
Instead of relying on isolated actions like form fills, agentic systems interpret multi-dimensional signals across accounts, buying group engagement, intent data, CRM history, and product usage, to understand when an account is truly in motion.

2. Agents
Agents turn signals into action. Each is designed for a specific objective, identifying in-market accounts, generating messaging, orchestrating campaigns, or surfacing sales insights, and improves over time based on outcomes.

3. Human oversight
Agentic GTM is not fully autonomous. Humans set strategy, define guardrails, and review high-impact outputs to ensure quality, compliance, and trust. Agents scale execution; humans ensure direction.

4. Orchestration
Agents must operate inside the systems where GTM teams already work, CRM, marketing automation, sales tools, and collaboration platforms, so insights translate into action in real time.

5. Outcomes
Every agent should be tied to a measurable business result: pipeline, deal velocity, win rates, or expansion. The shift is from optimizing activity to optimizing outcomes.

From funnel to system

Traditional GTM is linear:

Agentic GTM is continuous:

- Sense: Signals reveal what's happening across accounts
- Decide: Agents determine the next best action
- Act: Actions are executed across channels and teams
- Learn: Outcomes feed back into the system

The result is a GTM engine that is always on, continuously learning, and increasingly precise over time.

How to use this playbook

Each chapter maps to a layer of the model:

- Chapter 1 – Signals
- Chapter 2 – Agents
- Chapter 3 – Human Oversight
- Chapter 4 – Use Cases in Practice
- Chapter 5 – Trust and Orchestration
- Chapter 6 – Outcomes and Executive Activation

The goal isn't to adopt AI in pieces. It's to build a GTM system that works as a whole.

Ready? Let's get started.

01
From lead scoring to signal stacking

Featuring insights from Maximus Greenwald and Davis Potter

For years, B2B marketing and sales teams relied on lead scoring models to decide who to pursue. These models were often static, simplistic, and unreliable, based on arbitrary actions like downloading a whitepaper or opening an email. Too often, by the time a prospect had raised their hand by downloading a white paper, their mind was largely made up by all the research and evaluation they'd done before they ever raised their hand. The game has already started but sellers have started playing in the final minutes only.

Maximus Greenwald and Davis Potter propose something radically different: signal stacking.

Instead of chasing individuals based on one or two actions taken relatively late in the buying process, signal stacking uses AI agents to interpret multi-dimensional buying behavior across entire accounts throughout the buying journey. These agents combine CRM history, intent signals, buying group activity, email behavior, and product usage to deliver real-time visibility into which accounts are ready, and which aren't.

Maximus Greenwald
CEO,
Warmly,

- Advocates for replacing outdated lead scoring with real-time signal logic.
- His team uses AI to monitor group-level behavior patterns, not just individual activity.
- “Signal stacking is the new ICP. Combine buyer behavior, CRM, and historical outcomes to rank who’s ready to buy now.”

Davis Potter
Co-Founder,
Forge X

- Leads ABM systems built on pipeline-based signal tracking, not MQLs.
- Uses agentic AI to observe and orchestrate buying group engagement across the funnel.
- “You’re not scoring a lead. You’re modeling motion inside an account.”

Practical plays:

1. Use AI agents to track how many stakeholders are engaging within one account, mapping the buying group to inform engagement efforts.
2. Train models on closed-won opportunities to identify patterns of group engagement.
3. Replace lead scores with heatmaps showing total signal intensity across personas and channels.

This shift, from form fills to signal fields, makes the GTM engine smarter, faster, and more predictive. Instead of entering the game near the end, sellers can “sense and respond” to the moves of buyers from the start. That “gaining full visibility via signals, earlier in the game” approach is the foundation for every other agentic use case in this playbook.

02. Architecting agents with ROI in mind

Featuring insights from Chad Holdorf and Mark Boothe

It's easy to get distracted by shiny, newfangled tools. But real value from Agentic AI starts with asking and answering a foundational question: What is this agent supposed to do? Chad Holdorf and Mark Boothe stress that every AI agent, whether it's generating personalization, orchestrating campaigns, or filtering contact lists, must be built around a clear goal and measurable impact.

Chad Holdorf
VP Product,
Demandbase

- Built AI agents into Demandbase's platform using outcome-first logic.
- Advocates for filter agents that help non-technical teams interrogate massive datasets.
- "It's not about the tech, it's about the problem you're solving and the outcome you're driving."

Mark Boothe
CMO,
DOMO

- Created an Agent planning worksheet to define:
  1. **Agent objective**
  2. **Data required**
  3. **Workflow trigger**
  4. **Output format**
  5. **Success metric**
- "If you don't know why you're building it, you'll just build noise."

Practical plays:

1. Use Mark's worksheet before designing any AI use case.

2. Start pilots with low-risk, high-frequency tasks like email summaries or list segmentation.

3. Tie every agent to a pipeline, retention, or ACV goal, and measure results continuously.

Architecting AI agents with ROI in mind isn't just good strategy, it's the only way to ensure adoption, trust, and repeatable success.

03. The human-in-the-loop design principle

Featuring Myla Pilao and Chaz Horn

There's a temptation to treat Agentic AI like autopilot (i.e., set and forget). But the most successful teams don't aim for full autonomy. They aim for augmented intelligence, meaning AI that scales human creativity, decision-making, and empathy. To ensure both quality and scale, you want the best of human and artificial intelligence.

Both Myla Pilao and Chaz Horn stress that AI agents are most powerful when they're trained like teammates and treated like interns: fast, efficient, but still (and constantly) learning from others.

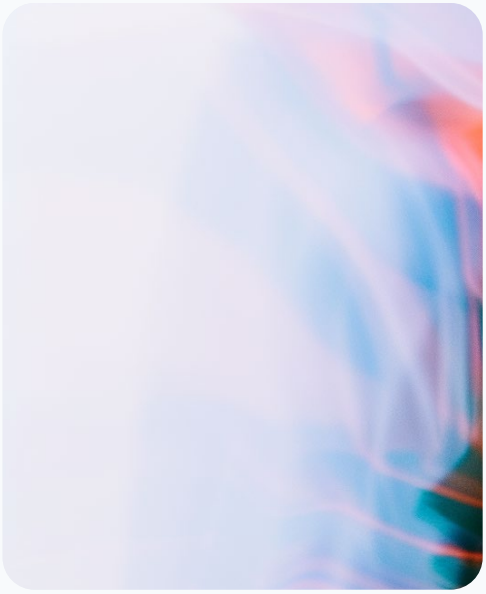

Myla Pilao
Dr. Technology Marketing,
Trend Micro

- Uses AI agents for foundational content creation, social copy, drafts, snippets.
- But humans always add the final tone, context, and compliance.
- “We pilot agents on low-risk content first… and iterate based on feedback.”

Chaz Horn
Founder,
Mastery Of B2B Sales

- Created a “clone” of himself as an AI agent, trained on his own frameworks.
- Clients can chat with the “AI Chaz” between coaching sessions.
- “They don’t get tired. They don’t react. But they still reflect your value system.”

Practical plays:

1. Use human-in-the-loop checkpoints for any public-facing or regulated content.
2. Create safe “playgrounds” inside your company to test and share prompt frameworks.
3. Assign human owners to each agent and create escalation logic when confidence is low.

Agentic AI isn’t about removing the human from the loop. It’s about amplifying the human touch at scale. With the right safeguards and feedback loops in place, AI agents can extend your brand’s voice, velocity, and vision, without compromising quality.

04.
ABM, content, and sales, use case spotlight

Featuring Vincent DeCastro, Matthew Creswick, and Shimon Ben Ayoun

While the potential of AI agents is massive, the value becomes tangible when they solve real problems for GTM teams. This chapter surfaces three high-impact use cases: (1) hyper-personalized ABM content, (2) always-on brand enforcement, and (3) real-time sales intelligence.

Vincent DeCastro
President,
The ABM Agency

- Uses AI agents inside the Manus platform to complete 5-person ABM plays in two days, not two weeks.
- Each agent performs a specialized task: research, message framing, content generation, layout, and localization.
- “You’re not just automating personalization, you’re scaling it across personas, languages, and stages.

Matthew Creswick
CMO,
Huble Digital

- Deploys AI agents to enforce brand governance and creative standards.
- Agents check content for tone, format, and message consistency across assets.
- “This isn’t about doing more. It’s about making everything tighter, faster, and on-brand.”

Shimon Ben Ayoun
Managing Director,
SPOTONVISION

- Empowers sales with agents that surface buying committee insights.
- AI helps sales understand who’s in the deal, what they care about, and how to position.
- “Sales needs more than contact info. They need context. Agents can provide that.”

Practical plays:

1. Use agents to generate personalized account decks, team cards, or value propositions.
2. Automate research and message creation for named accounts using Salesforce fields + AI prompts.
3. Integrate agents into Slack or CRM to deliver “Deal Briefings” on active pipeline accounts.

AI agents don’t just create content, they create precision at scale. With the right prompts, data, and oversight, they transform ABM from aspirational to operational.

05. Building operational trust through agent design

Featuring Jonathan Moran and Lukas Egger

It's not enough to launch agents, you have to trust them. And that trust must be designed into the system.

Jonathan Moran and Lukas Egger share a perspective rooted in systems thinking: AI agents must be explainable, composable, and governable. Otherwise, they erode confidence instead of scaling impact.

Jonathan Moran
Global Marketing Leader,
SAS

- Draws a sharp line between enterprise decisioning (governed, rules-based) and AI decisioning (adaptive, learning).
- Recommends combining both: structure + adaptability.
- "You can't scale decisioning if no one understands how or why it's working."

Lukas Egger
Head of Innovation
SAP Signavio

- Advocates for safe-to-fail environments with clear boundary conditions.
- Promotes experimentation over planning, deploy fast, fail fast, learn.
- "You don't plan your way into a revolution. You pilot your way in."

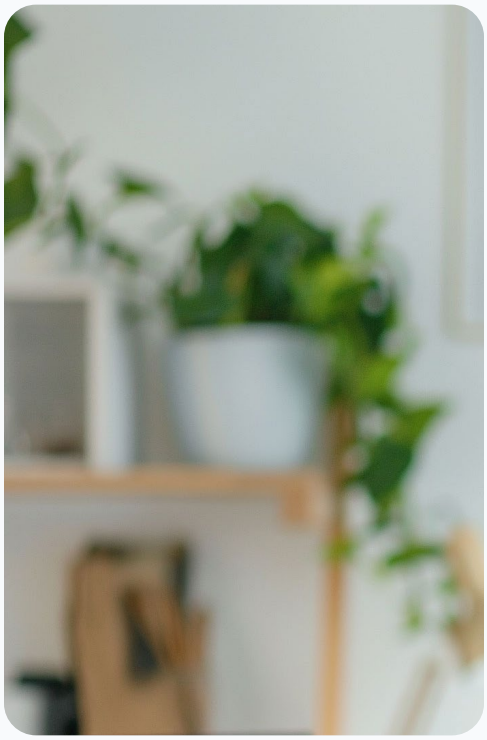

Practical plays:

1. Set “redlines” for agent actions, where human review is mandatory.

2. Embed agent decision trails into logs and dashboards.

3. Use modular architecture so agents can be turned on/off or swapped as needed.

Operational trust isn’t an output of Agentic AI. It’s a requirement for success. With the right mix of explainability, governance, and testing, teams can go from pilot to production with confidence.

06. Starting now, executive GTM activation guide

Featuring Daryn Smith, Scott Neuman, and Gabe Rogol

At some point, strategic thinking must give way to execution. This chapter is for GTM executives who are ready to move from awareness to activation.

What unites the most successful AI adopters isn't their tools, it's their mindset. They start small, know the specific goals they want to drive, move fast, and measure relentlessly. These leaders understand that Agentic AI is a long-term capability, not a short-term tactic.

Scott Neuman
VP Marketing,
Calix

- Uses agents to orchestrate actions across support, network ops, and marketing.
- Example: if a customer outage is detected, agents alert support, pause marketing, and update CRM.
- "AI helps us stop tone-deaf messaging before it happens."

Daryn Smith
CEO,
Huble Digital

- Recommends treating AI agents like interns, they learn by doing.
- Started with narrow pilots (e.g., proposal generation, internal Q&A agents).
- “Start where your team already works, plug into workflows they use daily.”

Gabe Rogol
CEO,
Demandbase

- Launched AgentBase as an orchestration layer for sales/ marketing teams.
- Focuses on shared pipeline goals, not leads.
- “Agents are now a GTM resource, just like SDRs or campaign managers.”

Executive activation checklist:

1. Identify one team (sales, marketing, CS) to run a pilot
2. Choose one use case tied to pipeline or revenue
3. Set up success metrics before launch
4. Assign a human owner to each agent or outcome
5. Review outputs weekly, adjust, and scale incrementally

Executives don’t need to master AI tech. But they do need to model AI-enabled leadership. That means embracing iteration, empowering cross-functional teams, and aligning AI output with business impact.

Agentic AI isn’t the future of GTM. It’s already here.
The question is: what will you build with it?

Appendix

15 Agentic AI use cases and prompts

This section offers tactical starting points for AI agent deployment, each use case paired with a prompt you can test, adapt, or scale today.

Sales enablement agents

1. Follow-up email rewriter
Prompt: “Rewrite this follow-up email to reflect the buyer’s role (VP Marketing) and reference our last demo discussion.”

2. Deal briefing generator
Prompt: “Summarize CRM activity, call transcripts, and recent web visits into a 3-point sales briefing for Account X.”

3. Objection handler assistant
Prompt: “Provide a 3-part rebuttal to a common pricing objection from a CFO in mid-market SaaS.”

ABM & campaign agents

4. Team card builder
Prompt: “Create a personalized team card for Acme Corp with role-specific messaging for Marketing, IT, and Procurement.”

5. Campaign variant generator
Prompt: “Adapt this 1-to-many campaign email into a 1-to-1 message for the VP of IT at a logistics firm.”

6. Industry personalization enhancer
Prompt: “Localize this content asset for a healthcare audience in the EU, referencing GDPR and data sovereignty.”

Research & intelligence agents

7. Buying group identifier
Prompt: “Analyze engagement data to infer roles involved in the last 10 closed-won deals and identify missing contacts in active pipeline accounts.”

8. Competitive summary generator
Prompt: “Summarize Competitor X’s Q2 earnings call highlights, positioning changes, and top product mentions.”

9. LinkedIn signal tracker
Prompt: “Identify job title changes, new hires, and content engagement trends for our top 25 target accounts.”

Content & messaging agents

10. Social post synthesizer
Prompt: “Create three LinkedIn posts based on this webinar transcript, one narrative, one stat-led, one opinionated.”

11. Buyer persona copy tuner
Prompt: “Rewrite this case study to resonate with a Head of Data Engineering in fintech.”

12. AI brand guardian
Prompt: “Check this landing page copy for brand tone alignment, overuse of buzzwords, and missing CTA.”

Operations & feedback agents

13. Agent performance analyzer
Prompt: “Show the win/loss impact and usage patterns for our Personalization Agent over the past 90 days.”

14. Campaign QA agent
Prompt: “Audit this nurture sequence for broken links, message redundancy, and deliverability flags.”

15. CRM clean-up assistant
Prompt: “Flag duplicate leads, missing firmographics, and outdated job titles in our top 100 accounts.”

Agent planning worksheet

Adapted from Mark Boothe's AI planning approach shared on the OnBase podcast

Before building or deploying any AI agent, use this worksheet to align your team on goals, data, workflows, and human oversight. This worksheet was inspired by the decision-making criteria and structured thinking emphasized by Mark Boothe.

Section	Description
Agent name	Give your agent a clear and functional name (e.g., "Onboarding Copy Agent").
Owner	Assign a human stakeholder responsible for the agent's output.
Objective	What is this agent meant to accomplish? (e.g., reduce onboarding email production time by 80%).
Primary users	Who will use or interact with this agent? (e.g., CS managers, SDRs, marketers).
Input data	What data sources does the agent need? (e.g., CRM, MAP, usage analytics, call transcripts).
Agent prompt / task	Write the core prompt or instruction the agent will follow.

Section	Description
Triggers	What event or schedule activates the agent? (e.g., new opportunity stage, weekly batch).
Output Format	What will the agent produce? (e.g., email draft, briefing summary, JSON block).
Delivery Channel	Where will the output be sent? (e.g., Slack, HubSpot, Google Doc).
Redlines / Human Review Points	What actions require manual approval? What can be fully automated?
Success Metrics	How will you measure performance? (e.g., time saved, engagement lift, error rate).
Feedback Loop	How will the agent be improved over time? (e.g., prompt iteration, user feedback form, monthly review).

Use this worksheet collaboratively across GTM, data, and enablement teams before deploying any production-level agent. It ensures that each agent is purposeful, measurable, and trusted.