



Webinar:

Early Wins in Agent AI for MarTech

How to pick the low-hanging fruit
(without slipping on a banana peel)

March 25, 2026

A Different Kind of Firm...

- MarTech Advisory and Consulting
- Deep understanding of the MarTech space
- An exclusive Council of top marketing experts across many verticals

Ask about RSG Membership and explore a seat on the private Council

GEEKS	MYSTICS
LOSERS	POSERS

“the mystical quadrant”

INDEPENDENCE FACTORS	REAL STORY GROUP	OTHER FIRMS
CONSULT TO VENDORS?	NO	YES
SPEAK AT VENDOR EVENTS?	NO	YES
ACCEPT EXPENSES & HOSPITALITY FROM VENDORS?	NO	YES
WRITE PAPERS FOR VENDORS?	NO	YES

Agenda



- 1 Where Agent AI in MarTech really stands today
- 2 Finding the early wins
- 3 How to not slip on the peel
- 4 The architectural question



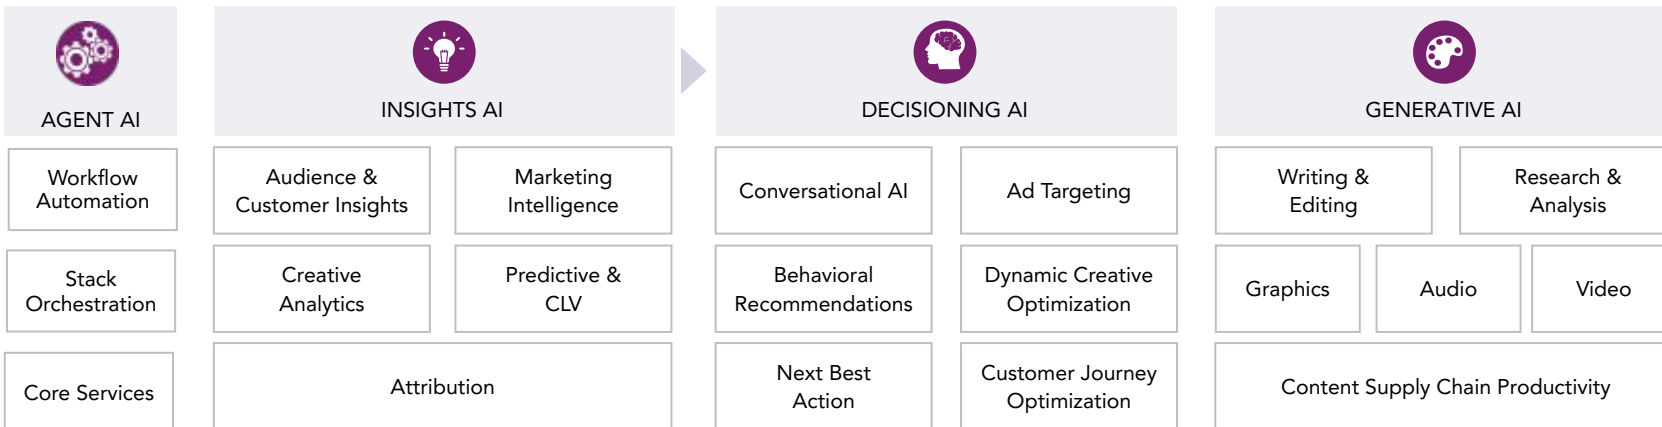
Where Agent AI in MarTech really stands today

Crawl-walk-run reality check and the agentic spectrum




MarTech AI Reference Model Overview

DIVERSE ENGAGEMENT SERVICE EXAMPLES



FOUNDATIONAL PLATFORMS

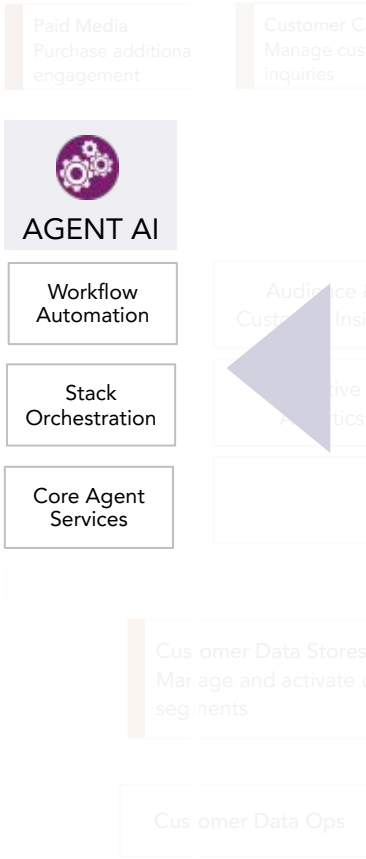


FOUNDATIONAL CAPABILITIES




RESPONSIBLE AI





1. What It Does:



Automates and sequences tasks within a platform (workflow) or across your MarTech stack (orchestration) — without needing human input at every step.


An independent, goal-driven system that uses data to plan, adapt, take action and improve.

2. Why It Matters:






By going beyond static decision-making, Agent AI introduces speed, autonomy, adaptability, and self-improvement to marketing operations.

Enables always-on, context-aware systems that potentially respond to change as it happens.



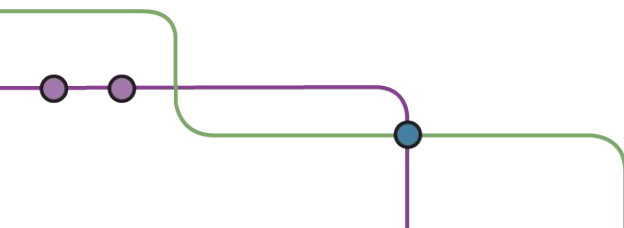
Crawl-Walk-Run Status for Agent AI Use Cases

Stage	Description	Reality for Most Enterprises
 Crawl	Early exploration or manual POC. Use cases are experimental, scripted, or HITL-heavy.	Where ~85-90% of RSG clients are today
 Walk	Repeatable workflows with some autonomy. Agents act partial orchestration or supervision.	10–15% are piloting here
 Run	Live, autonomous agents making goal-driven, cross-stack decisions in production.	RSG has not yet seen large enterprises undertake this, though many are exploring



Finding the early wins

Nine scenario categories, quick wins vs. strategic bets, and
how to score them



Agent AI Extends Established Use Cases & Offers New...

In most cases, agents extend existing use cases adding autonomy, memory, and goal-seeking to each siloed function (insight, decisioning, & generative)

New uses cases enable cross-functional workflows that link those capabilities into self-correcting loops

Traditional AI deployed tactically doesn't do this. Agent AI creates a "flywheel" effect of continuous self-improvement.



Nine agentic scenarios you can deploy now*

Insight Acceleration



Customer & Prospect Intelligence



Operational Analytics & Forecasting



Audience & Segment Optimization


Campaign & Execution Acceleration



Personalized Content Activation



Campaign & Execution




Orchestration & Decisioning


Support and Process Improvement



Content Workflow Automation



Support & CX Automation



Meeting & Interaction Intelligence

* To some degree: crawl/walk/run...

Example quick wins across channels

Email / Lifecycle: Dynamic blocks & subject lines with brand policy checks:

What it is

Agent assists writers by proposing subject lines and small content block variants for journeys.

Web / Commerce: Onsite search re-ranking & recommendations

What it is

Re-rank search results and nudge offers using behavioral signals and approved product content.

Ads / Acquisition: Guardrailed ad-creative variants & feed tuning

What it is

Generate/evaluate small creative or copy variants and optimize product/ad feeds for channels.

Content Ops: Content QA & auto-tagging in the supply chain

What it is

Speed review cycles by auto-tagging assets and flagging risk/consistency issues.

Contact RSG for details on...

- How it works
- Key success metrics
- Essential source alignment

Example strategic bets to queue next

Cross-stack Next Best Action & Journey Orchestration

What it is

Agents coordinate the next offer/action across site, email, paid, and service, not just within one tool.

Autonomous budget pacing & channel mix

What it is

Daily reallocation of spend across paid channels based on performance and constraints.

Suppression & rescue flows across the stack

What it is

Automatically suppress inapt audiences and trigger rescue offers across MAP, ads, and onsite based on CRM truth and recency.

Agent-assisted CX automation with human hand-off

What it is

Agents triage inbound (chat/case/email), propose responses, and escalate with a full decision trail.

Contact RSG for details on...

- How it works
- Key data and success metrics
- Essential source alignment
- Critical maturity and SDLC considerations

How to tell a quick win from a banana peel

Score each candidate scenario on five dimensions. High total = quick win. Low total = queue it or skip it.

Data readiness	Scope containment	Governance clarity	Failure tolerance	Measurability
<p>Quick win signal Clean, unified data already in a CDP or CRM</p>	<p>Quick win signal One or two systems; clear inputs and outputs</p>	<p>Quick win signal Approval rules and policies already documented</p>	<p>Quick win signal Low-stakes if agent gets it wrong; easy rollback</p>	<p>Quick win signal Clear KPI; easy before/after comparison</p>
<p>Banana peel signal Fragmented, stale, or no single customer view</p>	<p>Banana peel signal Spans 5+ tools; fuzzy handoffs</p>	<p>Banana peel signal No clear owner; compliance TBD</p>	<p>Banana peel signal Customer-facing, financial, or compliance-critical</p>	<p>Banana peel signal Vague ROI; hard to attribute outcomes</p>



How to not slip on the peel

Where pilots break, the four failure modes, and three fixes



Where Pilots Break

Data & meaning failure

What it looks like:
Agents act on stale or missing facts, cite outdated sources, or push the wrong audience/channel.

Why it happens: No CRM truth/identity, no recency rule, no approved source list—the agent lacks a common context.

Governance failure

What it looks like:
Content/flows move to calendars or live channels without claims/disclaimer checks; spend or volume spikes.

Why it happens: No approval pause and no budget/capacity caps before publish.

Observability failure

What it looks like:
You can't answer "why this audience/offer/budget?" for compliance or post-mortems.

Why it happens: No run log—no tools /sources /decisions /cost trail

Adaptivity failure

What it looks like:
Pilots never get smarter; repeat the same mistakes; can't learn from overrides.

Why it happens:
Outcomes and reason codes aren't fed back; no mechanism to pull last outcomes into planning.

Three solutions to the four problems

Context packet

A bundle of required facts the agent must read before it act.

Solves: Data & meaning (prevents stale/missing facts, wrong audience/channel) and enhances Adaptivity by pulling last outcomes into next run.

Approval gate & policy checks

A mandatory pause after drafts and before any calendar/publish.

Solves: Governance (prevents out-of-policy content and spend spikes).

Run log

A per-run decision trail

Solves: Observability (creates plan-vs-execution trail) and contributes to Adaptivity by storing reason codes/overrides for learning

Contact RSG for specific advice on these two...

Example coming up!



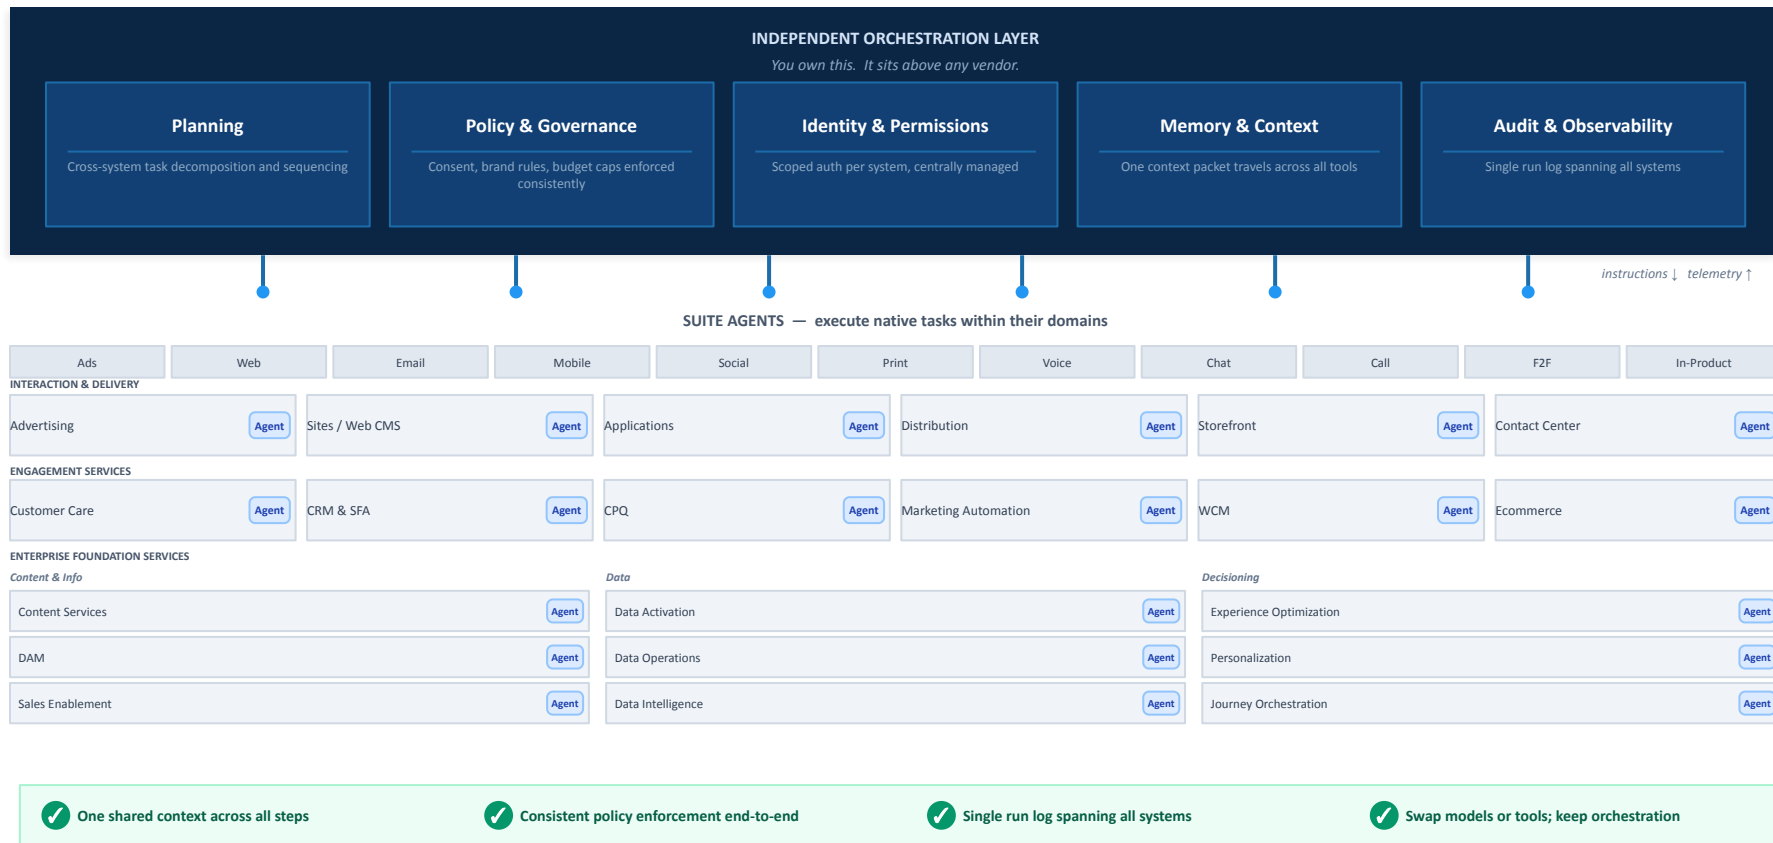
The architectural question

Suite agents vs. independent orchestration, and why it matters for scale



Independent orchestration, local execution

The control plane lives above the apps. Suite agents become workers, not orchestrators.



The architectural question you need to answer

Where does the control plane live?

Suite-first agents

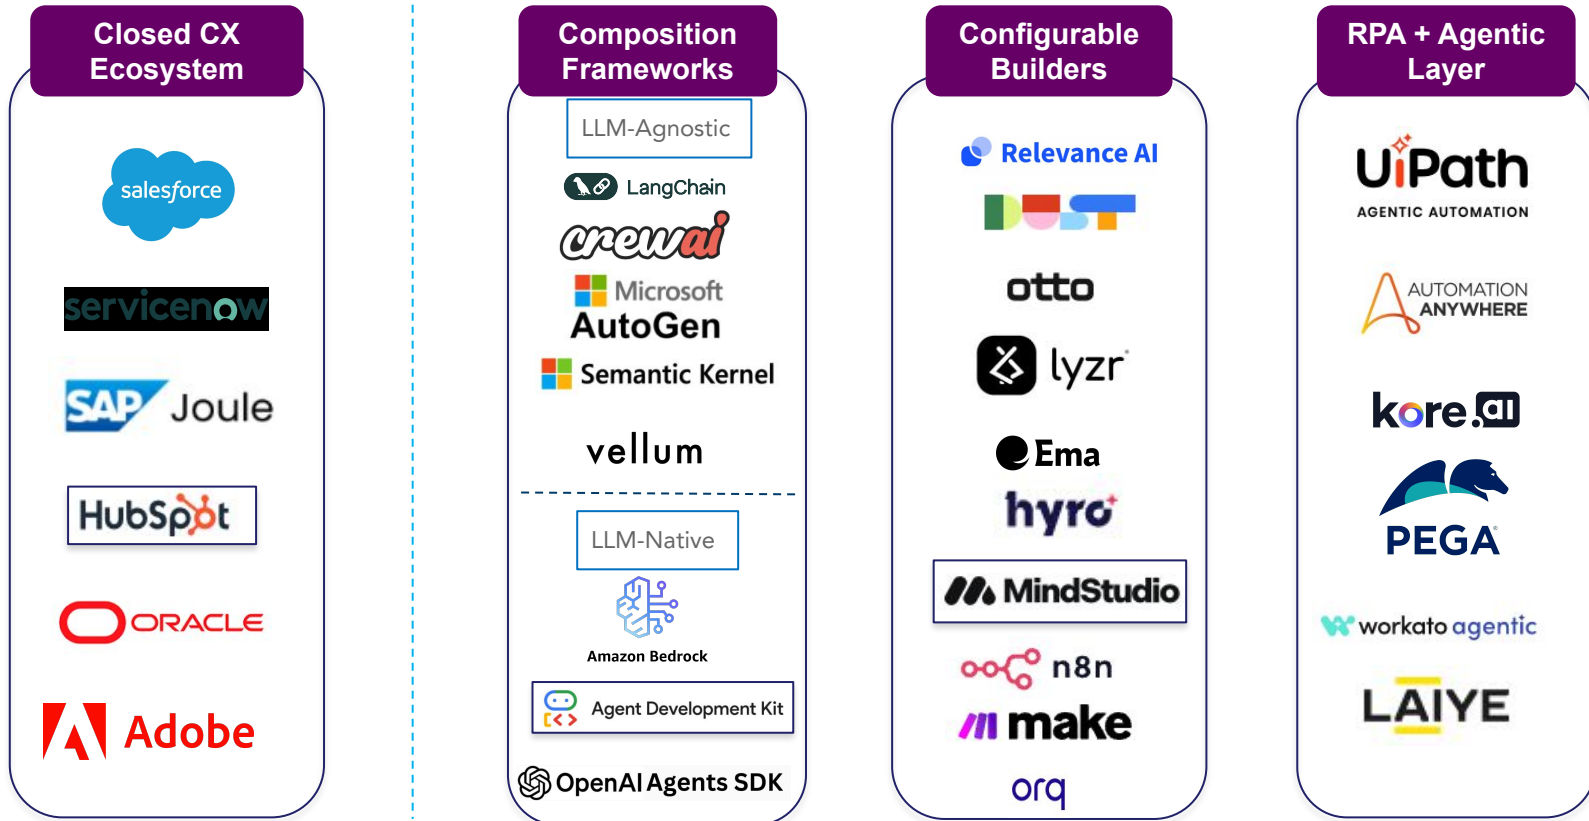
- Control plane lives inside the vendor platform
- Strongest behavior on the vendor's own objects and workflows
- Cross-silo reach routes through integration layers (MuleSoft, MCP, A2A)
- Each vendor optimizes locally within its walls
- Works well when the scenario lives inside one suite

Independent orchestration layer

- Control plane is yours to define (CrewAI, LangGraph, Pega, etc.)
- One context packet travels across all systems
- One run log spans everything for end-to-end audit
- Consistent policy enforcement across all vendors
- You own retries, timeouts, rollback, and auth scoping

If your scenario spans 3+ systems, needs consistent policy, or needs an end-to-end audit trail...
→ independent orchestration becomes the default.

Agent AI marketplace, Q1 2026



Agent Development Platforms

What MarTech Leaders Should Do...

Development Tactics

- **Business prototype, then pivot to code:** Let marketers sketch the flow; hand it to engineers before PII or money enters the picture
- **Budget for token (cost) audits:** Set a ceiling on spend. When the metric crosses that line, the workflow halts, alerts Ops, and prevents runaway bills or brand-damage
- **Mandate Red-Team reviews:** Run prompt-injection, data-leak, and adverse-output tests every sprint
- **Treat no-code like Excel macros:** Great for ad hoc tasks; reckless for core ops without IT oversight

Strategic Considerations

- **Balance value and speed:** For agentic use cases in particular, the highest-value use cases are also the most resource- / risk- / time-intensive; some quick-hit AI deployments deliver trivial returns
- **Vet in-platform AI features before lighting them up:** even simpler workflow enhancements may not fit within your enterprise guardrails; better to find out now rather than later
- **Don't underestimate the amount of real software development required to go from prototype to production (see left)**
- **Take a use-case approach to vendors and opportunities,** rather than deploy something ad-hoc just because it's available

How RSG can help

Turn pilots into reality with short, scenario-first engagements that add guardrails and produce a pragmatic plan.

GenAI for Content Supply Chain

- Assess real workflows across create, review, publish
- Test in-stack tools; add context, approval, run-log
- Deliver a scale plan with quality and governance gates

AI Use Case Rationalization

- Inventory AI features and use cases across teams
- Score by impact and feasibility; cut overlap
- Prioritize investments with a 4-quarter roadmap

Agent AI Readiness

- Identify where agentic flows make sense today
- Define context, HITL, and observability to ship safely
- Prove in a narrow slice before you scale

What you get

Content AI fit heatmap

Workflow evaluation results

Readiness checklist

Rationalization roadmap

Guardrails starter kit

Exec read-out with next steps

Optional services

AI Insights Accelerator

AI Scenario Modeling

AI-CDP Integration Review

GenAI Maturity Audit

AI Buying Decision Support

Composable AI Stack Design

Scenario-first. Stack-agnostic. Governance-backed. We test in your context and leave you with decisions you can act on.

Q&A

Stay in touch....



explore@realstorygroup.com

[linkedin.com/company/realstorygroup/](https://www.linkedin.com/company/realstorygroup/)

www.realstorygroup.com/Sample

Real Story
GROUP

Journey Orchestration Engine

Real Story
GROUP

Customer Data Platforms

Real Story
GROUP

Web Content & Experience Management

Real Story
GROUP

Digital & Marketing Asset Management

Real Story
GROUP

Email & Marketing Automation

Real Story
GROUP

Personalization Platforms

Real Story
GROUP

Omnichannel Content Platforms

Real Story
GROUP

Ecommerce Platforms
Comprehensive Product Evaluations

Real Story
GROUP

AI for Marketing
Comprehensive Product Evaluations