

The New Martech “Stack” for the AI Age

A Research Report
by Scott Brinker

In partnership with



Preface for the CMO	1
The Great Unstacking	3
The Arc of Martech Bends Toward Convergence	4
AI as the Catalyst for a 3rd Age	6
A Unified Data Foundation: Everything Is Data	9
The Convergence of Five Classes of Data	9
The Semantic Layer as the Keeper of Coherence	11
Context Graphs: A New Kind of Company Data	12
Convergence of Operational and Analytical Workloads	13
Universal Data Access: Data Without the Plumbing	14
More Software, Everywhere, By Everyone	16
The Software Palette of the Composable Canvas	16
The Hypertail: Custom Software at Scale	18
Governing the Explosion: Proliferation Without Chaos	19
Everything Assembles: Composability as the Default	20
The New Shape of the “Stack”	21
From Vertical Layers to Composable Capabilities	21
5 Rings of Capability, One Center of Gravity	22
Breaking the Integration Bottleneck	24
Composing Dynamic Partnerships	26
Getting There From Here	28
Start With the Data Foundation	28
Composability as Your North Star	29
From Boxes to Building Blocks	30
Conclusion	33

Preface for the CMO

Your martech stack has become both essential and exasperating. You’ve assembled dozens of tools over the years, each solving a real problem, but collectively creating a web of integration headaches. Now, AI is raising the stakes. New capabilities emerge weekly and your competitors are moving fast. Will your stack hold you back or help you pull ahead?

Marketing has the opportunity to lead enterprise AI adoption, as a full partner with your CIO or CTO. Your campaigns can be always-on, data-rich, and powered by continuous feedback loops — exactly what AI agents need to learn, decide, and act in real time. And because agents are only as good as the data they act on, they make your data foundation a strategic imperative, not just a hygiene issue.

This report proposes a fundamental shift in how marketing technology can be architected to give you the agility and power you need, transforming a rigid “stack” into a fluid **composable canvas** around a modern data platform such as Databricks. This transformation won’t happen overnight. But with a clear 3-5 year vision, you can start making decisions today that turn your stack from a constraint into a competitive advantage.

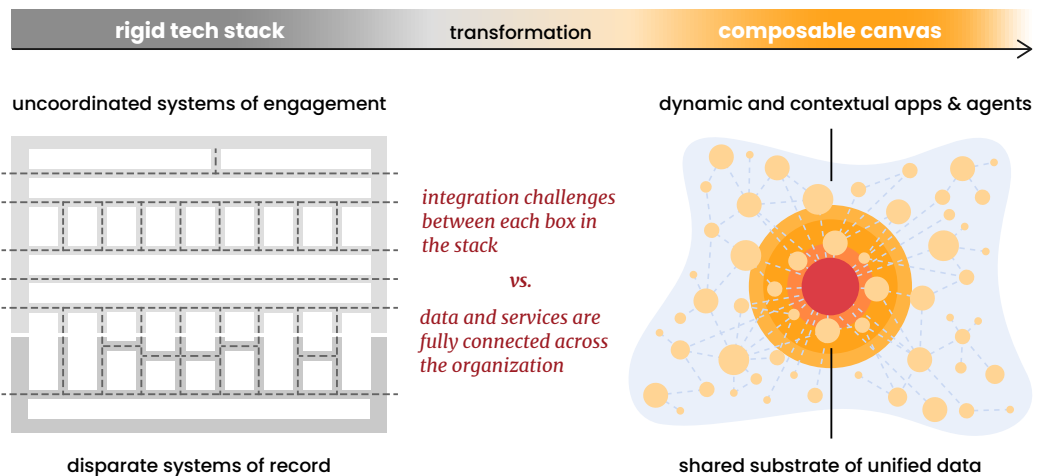


Figure 1.
The evolution of martech from a rigid stack to a composable canvas

What really matters, of course, is the impact this new architecture can have on your business outcomes:

Speed: From months to minutes. Today, adding a new marketing capability means integration projects, data pipelines, and IT tickets. In the composable canvas, new

tools and AI agents plug into a shared data foundation instantly. Spot an opportunity in a new channel? Want to spin up AI agents for a promising use case? You can move from idea to execution in days, not quarters.

Personalization at scale: Every customer, continuously. Your data lives in silos: customer profiles here, campaign performance there, content engagement somewhere else. Connecting who your customers are with what's actually driving results has been difficult to wrangle. The composable canvas unifies customer and company data into a single foundation where AI can orchestrate truly individualized experiences, informed by real-time performance, learning and improving with every interaction.

Reduced costs and complexity. Integration has been marketing's most expensive invisible tax. When systems share a common data foundation and speak common protocols, the integration burden drops by an order of magnitude. Those resources can shift from plumbing to programs that drive growth.

Adaptability as competitive advantage. Markets change. Customer expectations evolve. New channels emerge. A composable architecture lets you swap components, add capabilities, and reshape workflows without massive replatforming projects. Your technology becomes as agile as your strategy needs to be.

Differentiation through custom capabilities. When every competitor can buy the same martech products, the advantage goes to those who can tailor their own AI agents, specialized apps, and proprietary workflows that encode what makes your customer experience unique. The composable canvas makes this kind of custom development dramatically faster and more accessible.

“All CMOs are trying to solve the same three things: effectiveness, efficiency, and self-service.”

— Rick Schultz, CMO,
Databricks

This isn't a rip-and-replace proposition. The journey is incremental, and every step delivers value. But the companies that orient toward this architecture now will find themselves with marketing technology that adapts as fast as their markets demand.

This report draws on interviews with leaders who are already building toward this future: Meagen Eisenberg, CMO at Samsara; Bryce Peake, former VP of Marketing Decision Sciences at Domino's; Kumar Ram, VP/Global Head of Marketing Data Sciences at HP, and Chris Wissing, Chief Product Officer at Epsilon — as well as Databricks' own Rick Schultz, CMO, and Elizabeth Dobbs, AVP of Marketing Technology, Data & Growth. The pioneering work these leaders have already put into practice paints the way to the composable canvas ahead. Our deepest gratitude and thanks to them for their participation.

What follows shows how to get there. Your stack doesn't have to be a source of frustration. It can become a genuine competitive advantage and a huge unlock for your team's creativity and impact.

The Great Unstacking

The *martech stack*, once an obscure technical phrase, is firmly part of today's marketing lexicon.

Over the years, that stack has grown. It's not unusual for enterprise marketing teams to have dozens (and dozens and dozens) of products in their stack that they've accumulated. Its growth has been organic at most firms, in the same way that one might refer to an overgrown yard as "organic."

The upside of these best-of-breed stacks is they've been able to deliver innovative, specialized capabilities and add support for new and emerging channels and tactics. The downside is that integration — getting all these different tools to work together — has been a source of great frustration.¹

AI has put this all into hyperdrive. A whirlwind of new AI-native martech products. A torrent of new AI features added to incumbent platforms, which are fiercely determined not to let AI-native whippersnappers take their market share. And now a flood of custom-built AI agents, apps, and automations — yes, in some cases "vibe coded" — proliferating across the org.

Despite all these changes, the overall architecture of martech stacks has largely remained the same. So far. But it's creaking under its own weight, as new things keep getting thrown on top, stuck on the side, and wedged in between. The phrase "technical debt" is now becoming part of marketing's lexicon too. Every hour spent on integration work is an hour not spent on campaign optimization, personalization, or growth initiatives. The architecture of your stack has become a ceiling on marketing performance.

But what could the martech stack become in this new AI era?

That is the question this report proposes to answer. Or at least provide one possible answer. Because the architecture of the martech stack over the next 5-10 years doesn't have to look like it does today.

It's not just AI. Advances in cloud infrastructure by the hyperscalers and modern data platforms are enabling new ways to structure a company's overall tech stack and digital operations. The opportunity to collaborate across organizational boundaries with more partners and providers in a safe and seamless way is enabling new service-as-software vendor relationships and new ecosystem business strategies.

¹In the *AI & Data in Marketing Survey* run by chiefmartec and MartechTribe in Fall 2025, integration remained a Top 3 challenge for the majority of respondents.

"We've been trying to modify the same martech stack we've had since the internet started interneting. Folks, we're going to have to build a new one."

— Bryce Peake, former VP of Marketing Decision Sciences, Domino's

We have a generational opportunity to reimagine martech with this in mind.

We'll lay out a framework of a future martech stack architecture that we believe is ambitious but achievable. Most stacks will not look like this today, or even within the next year. But by gazing further out, on a time horizon of 3-5 years, we aim to set a North Star that can guide martech evolution in this new AI era.

Databricks sponsored this research and provided access to their platform, customers, and product team, giving us a laboratory to explore these concepts in depth. We use their platform as an example to illustrate our ideas. However, the concepts and principles here are intended to be portable and platform-agnostic.

Let's start with the big picture trajectory.

The Arc of Martech Bends Toward Convergence

We're at the beginning of a 3rd Age of Martech.

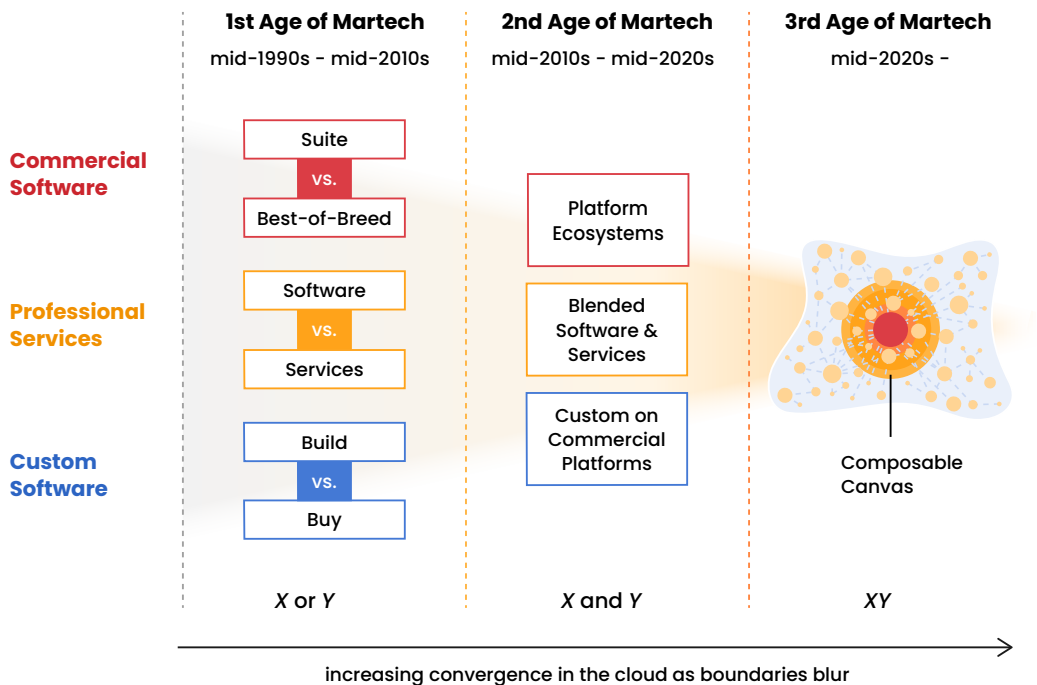


Figure 2.
The three ages of martech, progressively converging formerly disparate choices

The 1st Age of Martech (mid-1990s to mid-2010s) was defined by large, standalone products that eventually coalesced into multi-product suites. It was an age of dichotomies: suite vs. best-of-breed, software vs. services, buy vs. build.

But the rigidity of those dichotomies couldn't keep pace with marketing's accelerating expansion: a continuous flood of more data, more channels, more programs, more stakeholders, more governance requirements. The trade-offs between those either/or choices were painfully suboptimal.

The 2nd Age of Martech (mid-2010s to mid-2020s) began to break down those boundaries:

- From *suite vs. best-of-breed* to *integrated platform ecosystems*.
- From *software vs. services* to *blended delivery models with both*.
- From *build vs. buy* to *customization on top of commercial platforms*.

It's been an age of increasingly open systems. Products like customer data platforms (CDPs) and integration-platform-as-a-service (iPaaS) tools grew in popularity because they were built to serve a heterogeneous stack, moving data between systems and orchestrating workflows across individual products. Practitioners and vendors embraced the term "composable" to describe products and architectures designed for greater interoperability.

Yet the structure between these products remained rigid. Instead of a suite as one big box, martech stack diagrams became a mural of multiple smaller boxes, pieced together in a relatively fixed configuration, often not very well integrated:

- Multiple **systems of record** at the foundation: separate ones for customer data (CRM/CDP), content (CMS), brand assets (DAM), product information (PIM), etc.
- **Systems of engagement** layered above them: marketing automation platforms (MAP), digital experience platforms (DXP), e-commerce platforms, ad management platforms (DSP), business intelligence (BI), etc.²
- **Specialist apps** (and now **AI agents**) orbiting these systems of record and systems of engagement, for a wide range of use cases: webinars and event management, social media marketing, video/podcast production, customer profile enrichment, loyalty and referral programs, marketing attribution tools, and dozens of other niche applications.
- **Data platforms**, warehouses, lakes, and lakehouses, grew as a common substrate underneath the entire stack, pooling data from these different systems together. Initially this was for analytics, but the rise of "reverse ETL" and data clean rooms have made these data cloud platforms a more active part of marketing operations.

² The martech industry has generated TLAs — three-letter acronyms — faster than it's generated integrations between them.

The word “stack” was popularized because, well, the arrangement of these components are typically drawn as boxes, packed together in layers, stacked on top of each other.

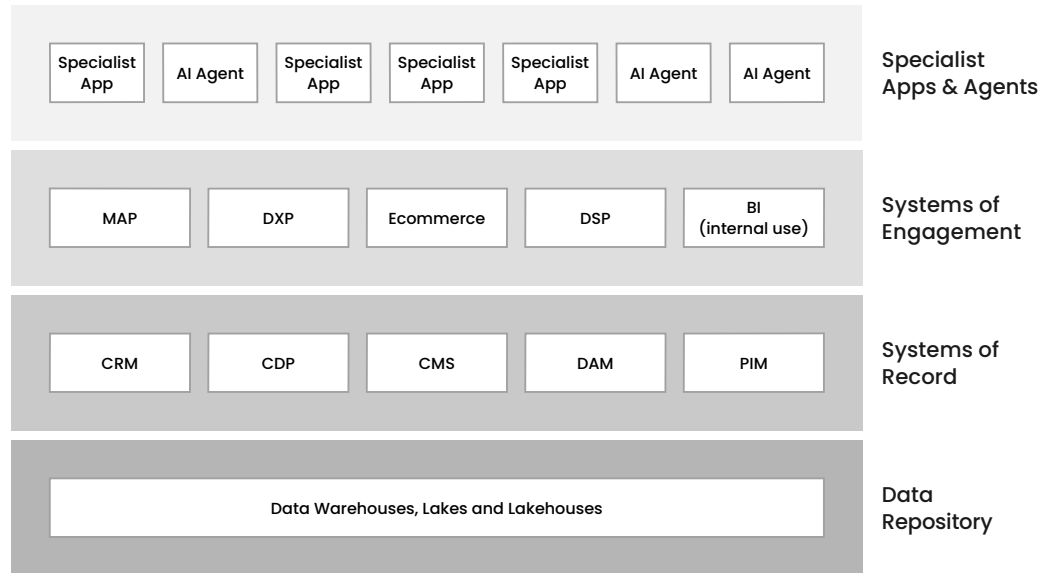


Figure 3.
The vertically layered structure of most martech stacks today, boxes on top of boxes

While more flexible than the monolithic suites of the 1st Age, this architecture has still been constrained. Integration between boxes is often limited to a subset of data and functionality. And systems beyond marketing — finance, product, IT, supply chain, HR — are rarely included.

True composability across the stack has remained aspirational.

AI as the Catalyst for a 3rd Age

AI is proving to be a catalyst for a more dramatic change.

Consider a B2B company with 10,000 accounts. Buying signals are firing constantly — website visits, content downloads, product usage spikes, support tickets, social mentions, job changes among contacts. Each signal is ambiguous on its own. Patterns emerge only when data from multiple systems are correlated in real time. And the “right” response depends on context. Is this a prospect or customer? What’s their contract status? Who owns the relationship? What actions have we already taken this week?

A dashboard can surface this information. But orchestrating the response — routing the right insight to the right person, triggering the right outreach, adjusting the right ads, updating the right nurture track — across thousands of accounts, continuously, in real time?

This is where dashboards end and agents begin.

The most valuable agentic workflows cross system boundaries as a matter of course. Agents are ravenous for “context” — relevant data, instructions, and tools for the specific goal they’re solving for. They crave data from across the organization, including behavioral streams and unstructured content that existed before but were too unwieldy to operationalize.

Simultaneously, a platform convergence is underway. Traditional SaaS vendors are racing to embed AI capabilities and agent frameworks into their products. Standalone AI agent builders are emerging as a new category, promising to orchestrate across any system. And data cloud platforms — the hyperscalers and modern data platforms — are pushing upward, arguing that since they already house the data, they’re the natural platform for software and AI operating with that data.

Everyone wants to be the place where agents run.

“The traditional layered structure can be better understood as a graph. It’s the foundation that consists of data and semantics, upon which a more fluid network of bi-directional communication between agents is built.”

— Tasso Argyros, VP Engineering, Databricks

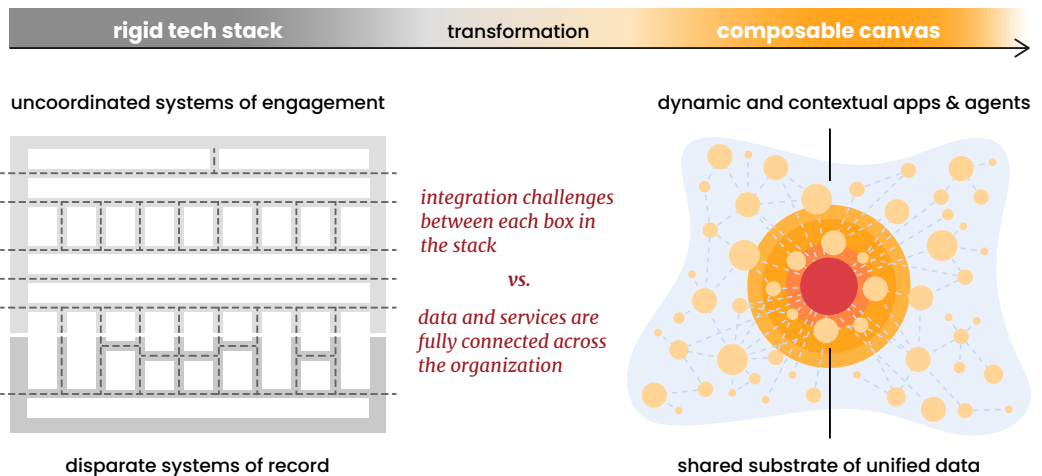


Figure 4. The evolution of martech from a rigid stack to a composable canvas

This pressure from above and below is compressing the stack. The neat layers of the 2nd Age are buckling. What emerges is better described as a unified **composable canvas**.

Agents and agentic workflows are, by their nature, composable. They assemble the data, tools, and capabilities they need dynamically, based on the task at hand. They

aren't anchored by hardwired, point-to-point integrations. They can adapt — if your architecture lets them.

The martech “stack” — that Tetris arrangement of boxes — starts to dissolve into a flexible workspace where data, content, and capabilities can be composed and recomposed on the fly. Rigid layers give way to a fluid architecture shaped by the work being done, not the products purchased.

This is the **3rd Age of Martech** and the beginning of The Great Unstacking.

A Unified Data Foundation: Everything Is Data

If the 3rd Age of Martech is defined by a composable canvas, then data is the canvas itself.

In the 2nd Age, data was something that moved between systems — extracted, transformed, loaded, synced, reverse-synced, and occasionally prayed over. Each product in the stack maintained its own data model, its own version of truth, its own gravitational pull. Integration was the art of building bridges between these data islands.

In the 3rd Age, data becomes the *shared substrate* on which everything else runs. Not data as an asset to be moved around, but data as a common foundation that agents, applications, and humans all operate upon together.

The martech stack doesn't sit on top of data. It is data.

We believe this is the **Grand Unifying Theory of Martech**: *everything is data*.

Customer profiles? Data. Marketing campaigns? Data. Brand guidelines? Data. The AI agent that writes your email copy? Also data — a model with weights, wrapped in prompts, governed by rules. Even the governance frameworks that keep everything in check are just another layer of data.

Once you see martech through this lens, the artificial boundaries between systems start to fall away. The question shifts from “how do I integrate these products?” to “what do I want my business operations and customer experiences to be?”

The Convergence of Five Classes of Data

What does “everything is data” actually mean in practice? Five classes of data — previously siloed in different systems, managed by different teams, and rarely brought together — now become interoperable on a unified foundation.

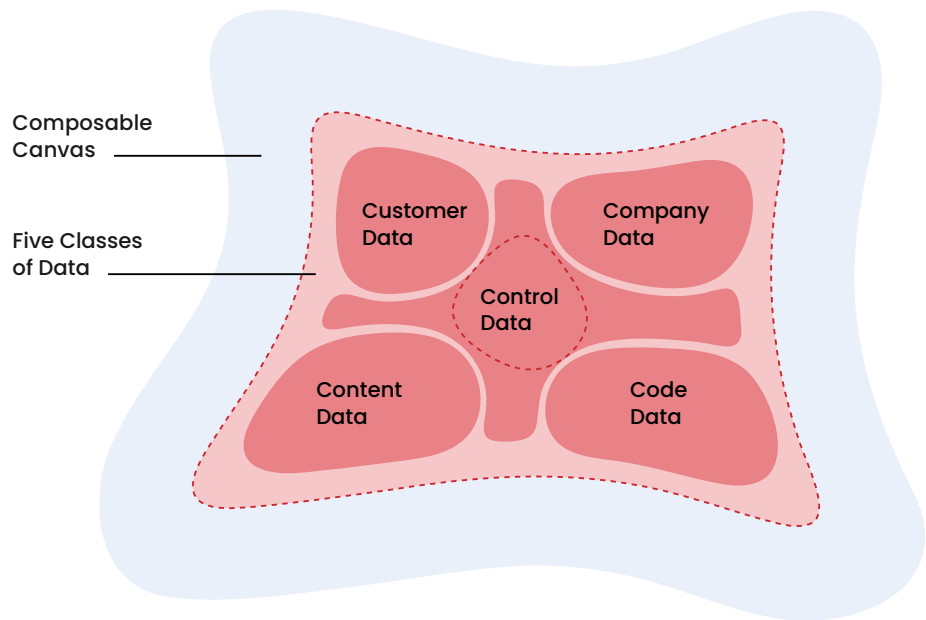


Figure 5.
Everything is data — five classes unified on one shared substrate

“People sometimes forget about how important campaign data is when talking about Customer Data Platforms (CDPs). I’m looking at campaigns all the time to figure out which ones we should run, which ones are working, which ones we should kill. If you can unify all your customer and campaign data — I sometimes call it the C-squared Data Platform — you can be smarter in your decisions.”

— Rick Schultz, CMO, Databricks

Customer Data is the class most familiar to marketers. This encompasses individual and account profiles, transaction histories, behavioral signals (web visits, product usage), emails, calls, chat transcripts, support tickets. It also includes second-party data shared by partners and third-party data enriched from external providers. In the 2nd Age, this was the domain of CRMs and CDPs. In the 3rd Age, it remains central but is no longer locked inside any single system.

Company Data is operational data from across the organization, not just marketing: inventory levels, financial records, campaign spend, sales pipeline, logistics, and supply chain status. This is the data that marketing has historically had to try to wrangle out of point solutions or request from other departments through formal processes (or informal favors). When it lives on a shared foundation, marketing gains context that was previously invisible — and other departments gain visibility into marketing’s activities in return.

This visibility changes how marketing proves its value. When campaign data connects to pipeline and revenue data in the same foundation, attribution stops being a quarterly debate and becomes a continuous, fully aligned conversation with the CFO.

Content Data reframes creative assets as structured data. Brand guidelines become queryable. Product information, images, videos, and copy are stored with rich metadata that describes not just what they are, but where they can be used, who approved them, and how they perform. This is more than a traditional DAM with better tagging. It’s content treated as a first-class data citizen, ready to be assembled and personalized dynamically by agents that understand the rules.

³ “In computer science, the expression code as data refers to the idea that source code written in a programming language can be manipulated as data.” – Wikipedia, [Code as data](#)

Code Data may be the most unfamiliar to marketing, but it is increasingly essential to their operations: *software is data*.³ AI models and their weights. Prompts and prompt templates. Agent skill definitions. Source code for custom applications. Configurations for agents and automations. They’re artifacts that can be versioned, governed, and orchestrated like any other data asset. When your marketing automation is a prompt chain rather than a flowchart, the prompt is the logic.

Control Data is the connective tissue that makes the other four classes coherent and compliant: semantic layer definitions that create shared meaning across systems, governance policies that determine who can access what, business rules that encode organizational logic, and guardrails that keep AI agents operating within acceptable bounds. Control Data transforms a data lake from a swamp into a navigable waterway.

The Semantic Layer as the Keeper of Coherence

When everything is data, the risk is that everything becomes a mess.

Different systems use different terms for the same concepts. What sales calls an “opportunity” might be what marketing calls a “qualified lead” and what finance calls “projected revenue.” Customer lifetime value can have a dozen definitions within a single organization, depending on who’s calculating it and what they’re trying to prove.

This is where a **semantic layer** becomes essential. A semantic layer sits above the raw data and provides a consistent, business-friendly vocabulary for querying it. It defines what “customer” means across the organization. It standardizes metrics so that everyone is working from the same calculations. It translates between the technical schemas of a myriad of data sources and the conceptual models that business users actually think in.

The semantic layer takes on heightened importance in the 3rd Age because AI agents rely on it. An agent trying to identify high-value accounts at risk needs a shared definition of “high-value” and “at risk.” Without a semantic layer, every agent becomes its own island of interpretation — which is how you get three different dashboards showing three different pipeline numbers, or worse, three different agents taking three different actions based on contradictory assumptions.

The semantic layer is the keeper of coherence. It doesn’t just make data accessible; it makes data *meaningful*.

“It used to be that content would be the last consideration in the marketing process; however, with agents taking the lead, and brands codifying their brand guidelines as agentic context, content becomes the tip of the spear for true personalization; it moves from being the bottleneck at the end of the process to the main leverage point that agents have to drive conversion.”

— Tasso Argyros,
VP Engineering, Databricks

Context Graphs: A New Kind of Company Data

Traditional systems of record capture *what* happened. What they don't capture is *why* it was allowed to happen.

The reasoning behind decisions — the exceptions granted, the precedents invoked, the cross-system context synthesized, the approvals obtained — currently lives in ephemeral places: Slack threads, Zoom calls, email chains, and people's heads.

A VP approves a 25% discount on a renewal call. The CRM shows the number; it doesn't show who approved the deviation, what service incidents justified it, or which prior deal established the precedent. And when an AI agent makes a recommendation — or a human overrides one — that rationale is equally ephemeral.

A new concept called a **context graph** addresses this gap. Yes, "context" is doing a lot of work these days — from context windows in LLMs to context engineering for agents to the everyday sense of "reading the room." But as recently articulated by Jaya Gupta and Ashu Garg of Foundation Capital⁴, a context graph means something more specific: a living record of decision traces, stitched across entities and time, that makes precedents searchable — whether those decisions were made by humans, agents, or both.

For marketing, this has immediate practical implications. When a lead gets routed to sales, the context graph records why: which signals fired, which scoring model was applied, what threshold was met. When a campaign budget gets reallocated mid-quarter, the graph captures the approval chain and the performance data that justified it. When an agent proposes a personalized offer and a human approves it, both the recommendation logic and the approval become durable artifacts.

This creates a feedback loop that compounds over time. Captured decision traces become data that AI agents and human analysts can reason over. Similar situations can reference how they were handled before. Exceptions don't have to be re-litigated in Slack every quarter. They're documented, queryable, and available to inform both human judgment and agent behavior.

The context graph represents a new type of Company Data that captures organizational intelligence never systematically stored before. Some have argued that context graphs require their own dedicated layer in the tech stack, separate from the data foundation. We see it differently: decision traces, reasoning chains, and approval histories are *data* — they can be structured, governed, versioned, and queried just like customer records or financial transactions. The most natural place for that to take place is within the unified data platform, where the rest of the enterprise data lives.

⁴AI's trillion-dollar opportunity: Context graphs, Jaya Gupta and Ashu Garg, December 2025

"We think of marketing data in terms of what decision we're trying to drive. Rather than mirroring the business's org chart, we organize our semantic layer by decisions that span multiple teams or functions. It's the best mental model for organizing data, and it makes data easier to find. When you're mapping a customer journey, you're looking across specific decision spaces: consideration, deal consideration, value consideration. Decision data becomes intuitive to navigate."

— Bryce Peake, former VP of Marketing Decision Sciences, Domino's

Treating context graphs as data isn't just philosophically consistent with our "everything is data" thesis. It's practically essential. When the context of decisions lives in the same foundation as the other data it references, agents can seamlessly connect the "what happened" with the "why it happened." Separate the two, and you've recreated the very silos the composable canvas is designed to eliminate. The context graph may prove to be the most strategically valuable addition to the data foundation in this new era.

Convergence of Operational and Analytical Workloads

For decades, operational systems (where transactions happen) and analytical systems (where data is analyzed) were kept separate. The conventional wisdom held that you couldn't run real-time operations and complex analytics on the same infrastructure without one degrading the other.

Modern data platforms are collapsing this divide. They can handle both the transactional demands of operational workloads and the computational intensity of analytical queries, often on the same underlying data.

For martech, this convergence is transformative. It means the data that powers your dashboards can also power your agents, without extraction delays or synchronization headaches. Real-time behavioral streams can flow into the same environment where historical analysis is performed. An agent can query customer lifetime value, check current inventory, and trigger a personalized offer, all from one unified foundation.

An example in practice: this is the difference between a loyalty program that reacts to last month's purchase history and one that responds to what's happening right now — adjusting offers based on real-time inventory, current margin targets, and the customer's behavior in this session.

This doesn't eliminate the need for specialized applications. Your e-commerce platform still needs to process transactions at scale. Your marketing automation still needs to send emails reliably. But these systems increasingly write to and read from a common data layer, rather than maintaining their own isolated data stores.

The convergence of operational and analytical workloads is what makes "everything is data" more than a philosophical position. It's the technical capability that enables it.

Universal Data Access: Data Without the Plumbing

For most marketers, data integration has always meant waiting. Waiting for IT to build a connector. Waiting for a sync to complete. Waiting to find out why the numbers in one system don't match the numbers in another. The plumbing required to move data between systems has been one of the great time sinks of modern marketing operations.

The unified data foundation of a composable canvas changes this equation. Instead of data living in isolated systems that must be laboriously connected, data becomes universally accessible — queryable from wherever it's needed, regardless of where it physically resides.

This shows up in three increasingly common patterns:

“Once all of our marketing data was centralized in the lakehouse, we revisited models that had never fully leveraged that richness and asked: how much performance lift are we leaving on the table? We started with lead scoring, partly because our external vendor was costly and scaling data ingestion into that system wasn't sustainable. The impact was immediate — the rebuilt model converted at 4X the previous rate. From there, we quickly replicated the same approach to launch a highly predictive account scoring model in just four weeks. It fundamentally changed how we think about building on our own data foundation.”

— Elizabeth Dobbs, AVP of Marketing Technology, Data & Growth, Databricks

Applications built on the data layer. A new generation of martech tools — including many composable CDPs — don't maintain their own separate databases. They operate directly on data in the warehouse or lakehouse. When you build an audience segment, you're not querying a copy of your customer data that was synced overnight; you're querying the actual customer data. There's no integration to build because the application and the data already share the same foundation.

Queries that cross system boundaries. Modern data platforms can reach into external systems — other databases, cloud applications, legacy systems — and query them as if the data were local. From a marketer's perspective, this means that customer records in the CRM, behavioral events in the product analytics platform, and transaction history in the ERP can all be combined in a single analysis or activation, without anyone having to build a pipeline to make that possible.

Data sharing across organizations. When you want to enrich your customer profiles with data from a retail partner or a third-party provider, the traditional approach meant negotiating file transfers, building import processes, and reconciling schemas. Modern data platforms enable something more elegant: the data partner can grant you access to query their data directly, governed by policies they control, without either party having to move or copy anything.

Data engineers describe these capabilities using terms like “query federation” and “data virtualization.” In the martech world, you'll often hear this referred to as **zero-copy** architecture, a nod to the fact that data doesn't have to be duplicated into yet another system to be useful. (It's worth noting that “zero-copy” has become something of a marketing term itself, and different vendors define it differently.

“That openness at the storage layer gets rid of the locked-in aspect that’s plagued this industry. We wanted something that was portable and open so that we weren’t locked into a particular cloud.”

— Chris Wissing, Chief Product Officer, Epsilon

The core idea — reducing or eliminating redundant data movement and redundant governance overhead — is sound, even if the implementations vary.)

The practical implication: the data you need is increasingly *just there*. Not after a sync. Not after a ticket to IT. Not after reconciling three different versions of truth. For marketing, this translates directly to speed: campaigns that launch in days instead of weeks, audiences built in minutes instead of hours, and personalization informed by today’s behavior rather than last month’s data extract.

This is what enables the composable canvas. When agents, applications, and analysts can all access the same unified data foundation, regardless of where individual data assets physically reside, the rigid boundaries between systems start to dissolve. The architecture becomes fluid, shaped by what you’re trying to accomplish rather than by which vendor happens to store which data where.

More Software, Everywhere, By Everyone

If the unified data foundation is the canvas, software is what gets painted on it.

And there's about to be a lot more painting.⁵

The martech landscape has grown from roughly 150 solutions in 2011 to over 15,000 today.⁶ That growth was driven primarily by commercial software, vendors building products to sell to marketers. But that commercial expansion, as dramatic as it's been, is only the prelude. The next wave of growth won't come mainly from vendors. It will come from within.

AI has fundamentally changed who can create software and how fast they can create it. Professional developers using AI coding assistants are dramatically more productive. Marketing ops professionals and other "citizen developers" can now build functional applications through natural language prompts — the phenomenon sometimes called "vibe coding." And perhaps most consequentially, AI agents themselves can generate software on-the-fly to accomplish specific tasks, then discard it when the job is done.

This changes the math of what's possible. When building an app takes months and costs six figures, you only build what's absolutely necessary. When building an app takes hours — or minutes, or seconds — you build whatever would be useful. Ideas that never stood a chance before can now be tested by Tuesday.

The Software Palette of the Composable Canvas

Three types of software operate on this canvas:

Platforms are the foundational systems that enable services, apps, and agents to share data, coordinate actions, and scale together. Primarily anchored by major players — hyperscalers (AWS, Microsoft Azure, Google Cloud), modern data platforms (Databricks, Snowflake), and core business systems (Salesforce, SAP, Workday) — they're almost exclusively commercial. Few companies build their own foundational platforms from scratch.

⁵ Data isn't the new oil. It's the new oil paint.

⁶ *2025 Marketing Technology Landscape Supergraphic: 100X growth since 2011, but now with AI...*, chiefmartec, based on research by chiefmartec and MartechTribe, May 2025

"Imagine a UX designer who monitored how you use an app every day, then customized it for your specific needs overnight. That's where we're heading — software that adapts to your workflow rather than forcing you to adapt to it."

— Tasso Argyros, VP Engineering, Databricks

Services are reusable capabilities that can be composed into larger solutions. Some are API-first by design — payment processing (Stripe, Adyen), email delivery (SendGrid, Postmark), content delivery (Contentful, Storyblok), data enrichment (Cognism, ZoomInfo). Others are horizontal and general-purpose — most notably foundation model providers (OpenAI, Anthropic, Google), whose APIs have become foundational building blocks for the entire canvas. What they share: on-demand access to capabilities through programmatic interfaces.

Services are expanding in two directions.

First, established martech applications are increasingly exposing their functionality via APIs, not just their data. A marketing automation platform isn't just a place to build campaigns through its UI. It's a service whose campaign execution capabilities can be invoked by agents and custom applications. The UI remains valuable for human users; the API makes the same capabilities available to software.

Second, the boundary between software and services is dissolving from both sides.

On the provider side, *software is eating services*. AI enables service providers to operate at a speed, scale, and cost structure that transforms their value proposition: localization completed in hours instead of weeks, competitive intelligence updated continuously rather than quarterly, creative production at 10x the volume and a fraction of the timeline. This is what the market increasingly calls **service-as-software**.

On the client side, *services are becoming software-accessible*. Agencies, research firms, and specialized consultancies are exposing their capabilities through programmatic interfaces. Creative assets delivered directly into the DAM. Research findings structured for ingestion into analytics platforms. Briefs submitted and deliverables received via API. What happens inside the provider — human-intensive, AI-assisted, or fully automated — is increasingly invisible. The service becomes something that agents and applications can invoke as part of automated workflows.

Apps and Agents are where the explosion really happens. Some will be commercial products — the “long tail” of specialized martech solutions that address specific use cases. These tend to be applications where the user interface is central to the value (creative design tools, video editing platforms), where the capability is broadly applicable across many companies (webinar platforms, survey tools, event management systems), or where the technical sophistication requires dedicated R&D investment.

But alongside this commercial long tail, a larger **hypertail** of custom software is emerging — applications and agents built by companies for their own specific needs. If a company might purchase dozens or hundreds of commercial apps, it might eventually operate thousands of custom ones.

“My view on build vs. buy has really changed. If you had asked me a year ago, I would have talked about a set of off-the-shelf SaaS tools. Now it's all about what we can do with AI. We have 25 AI-native tools we're immersed in and training our marketers on, plus 12 agentic AI use cases we've created that are unique to our business.”

— Meagen Eisenberg, CMO,
Samsara

The Hypertail: Custom Software at Scale

The hypertail encompasses several types of custom software, distinguished by who builds them and how long they last.

IT-built applications are the traditional model of custom development: internal teams building bespoke solutions for specific business needs. AI dramatically accelerates this work, but the basic pattern is familiar: a defined project, a development process, a deployed application maintained over time.

Citizen-developed applications emerge when marketing ops professionals, revenue operations teams, or individual contributors use AI-assisted tools to build their own solutions. Workflow automations, custom dashboards, specialized calculators, simple apps that solve specific problems. The person who understands the need is empowered to address it directly, without filing a ticket.

Agent-generated software is the most novel category — code that AI agents create on-the-fly to accomplish specific tasks. An agent orchestrating a complex campaign might generate a custom script to transform data between systems, use it once, and discard it. Another might create a micro-application to handle a specific customer inquiry, then dissolve it when the interaction ends.

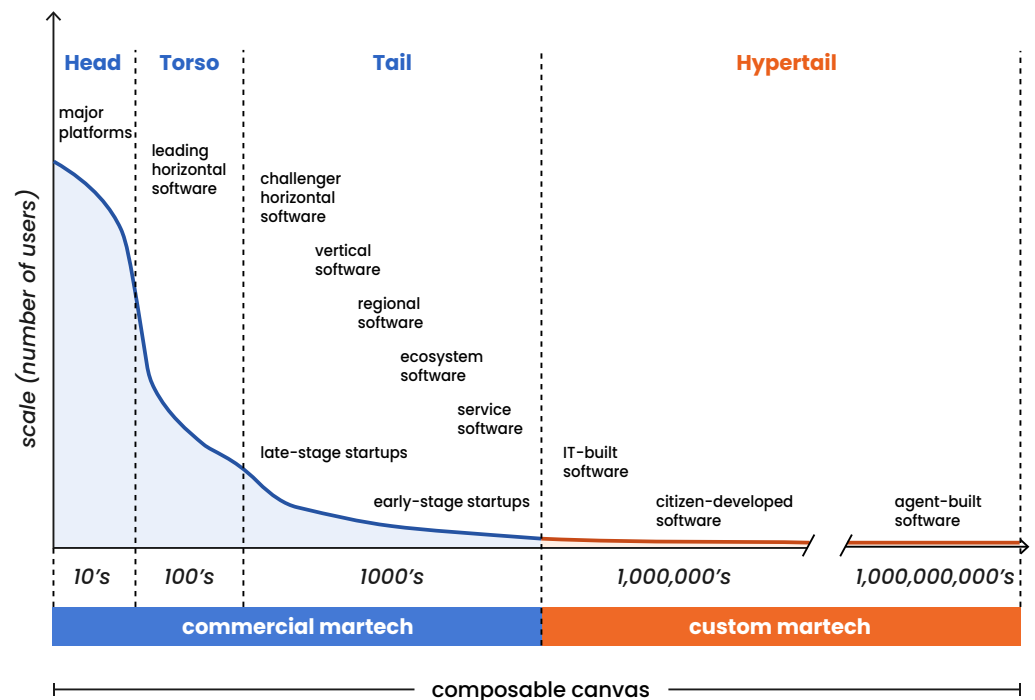


Figure 6.
The “long tail” of commercial martech and the “hypertail” of custom-built martech

This last category challenges traditional notions of what “software” means. It can be ephemeral, rather than persistent. Generated rather than developed. Just-in-time rather than planned. Call it **generative software** — code that exists briefly to serve a specific moment, then vanishes.

The hypertail is where differentiation lives. Commercial software is built to be sold to many companies. Your competitors can buy the same products you can. Custom software, in contrast, captures what makes your company unique.

This is where your proprietary understanding of customers becomes encoded in systems that act on it: a B2B company’s lead scoring model tuned to its specific ICP signals, a retailer’s personalization engine trained on its unique product relationships and customer segments. The advantage isn’t the data alone — it’s the tailored functionality you build on top. The hypertail is where you build what competitors can’t buy.

Governing the Explosion: Proliferation Without Chaos

More software raises the stakes for rock-solid, consistent governance.

“If you’re familiar with the whole ‘GEO is cannibalizing SEO’ trend — our team actually built a tool for this in-house. They created a Streamlit app that sits on top of Databricks and runs over 400 prompts twice a day, learning what prospects and customers are searching for and where we’re getting citations. We’re able to rapidly adapt and optimize our content based on these insights.”

— Meagen Eisenberg, CMO,
Samsara

When custom applications can be spun up by anyone with a prompt, how do you ensure they’re using data appropriately? When agents generate ephemeral code, how do you audit what happened? When every team can create their own solutions, how do you maintain coherence?

This is where earlier investments in the unified data foundation pay dividends. The semantic layer doesn’t just provide shared meaning for data. It provides shared definitions for operations. If “customer lifetime value” means the same thing regardless of which app or agent calculates it, you get consistency without centralized control. If access policies are enforced at the data layer, you don’t have to trust every citizen-developed app to implement its own security.

The same principle extends to code itself. When we treat code as data, the semantic layer can govern custom apps and agents just as it governs other data definitions. It can ensure that common operations — sending a message, calculating a discount, accessing a customer record — execute consistently regardless of which app or agent invokes them. Version control, access policies, and audit trails keep the system transparent, compliant, and secure.

“We’re seeing the emergence of app-level point solutions as a real alternative to traditional SaaS purchases. The more effectively we ingest and manage our data in our marketing lakehouse, the more powerful and effective those applications become.”

— Elizabeth Dobbs, AVP of Marketing Technology, Data & Growth, Databricks

The result is proliferation without chaos: thousands of custom apps and agents, all operating with shared meaning.

There’s a continuum between consistency and customization that every organization must navigate. Some processes benefit from standardization. You probably want everyone calculating pipeline the same way. Others benefit from local optimization. A regional team might need workflows tuned to local market dynamics. The art is knowing which is which, and having an architecture that supports both.

Everything Assembles: Composability as the Default

The through-line across all three types of software — platforms, services, apps and agents — is true *composability*.

Platforms expose capabilities through APIs. Services are designed to be assembled. Apps and agents are built from components. Even the hypertail of custom software is constructed by composing existing building blocks rather than starting from scratch.

This composable architecture favors open approaches. Open data formats, such as Linux Foundation Delta Lake⁷ and Apache Iceberg⁸, prevent platform lock-in. Open protocols, such as the Model Context Protocol⁹ for agents, enable interoperability across tools. Open source components provide building blocks that anyone can inspect, modify, and extend.

The 2nd Age saw a shift from closed suites to more open ecosystems. The 3rd Age pushes this further. Composability becomes the default assumption, not a feature to be negotiated. The question isn’t whether systems will connect, but how fluidly they can be assembled and reassembled as needs evolve.

In a composable world, the boundaries between “bought” and “built” blur. A solution might combine a commercial platform, several commercial and custom services, and dozens of custom apps and agents — all operating on a unified data foundation, all governed by consistent policies, all part of the same canvas.

⁷ [Delta Lake](#), The Linux Foundation Projects

⁸ [Apache Iceberg](#), The Apache Software Foundation

⁹ [What is the Model Context Protocol \(MCP\)?](#), The Linux Foundation Projects

The New Shape of the “Stack”

For two decades, we’ve talked about martech “stacks.” The metaphor made sense: layers of technology built on top of each other. Data at the bottom. Applications in the middle. Experiences at the top. Neat, orderly, vertical.

The problem is that stacks are rigid. You can’t easily remove something from the middle without destabilizing what’s above. You can’t rearrange the layers. Everything has a fixed position, and change is expensive.

This made some sense when martech was mostly about buying software and bolting it together. The stack metaphor reflected reality: big platforms at the foundation, specialized applications layered above, integrations connecting the pieces. Change was hard because change was hard — technically, contractually, organizationally.

But that reality is shifting. When software can be generated on-demand, when data flows through a unified foundation, when capabilities can be composed dynamically, the vertical stack metaphor stops serving us. It constrains our imagination.

The 3rd Age needs a different shape.

From Vertical Layers to Composable Capabilities

The shift is from thinking in layers to thinking in *capabilities*.

In the old model, you asked: “What layer does this belong to? What sits above it? What sits below?” Architecture was defined by position.

In the new model, you ask: “What capabilities do I need? How do they connect? How easily can I swap one for another?” Architecture is defined by composition.

This is more than semantics. It’s a fundamental change in how martech systems are designed.

Logical design over “physical” design. In the old model, architecture mapped to distinct, bounded applications — solid boxes, each a self-contained bundle of data, logic, and interface. Changing what your martech did meant replacing these boxes wholesale. In the new model, capabilities compose on a shared substrate. You reshape the logic without having to move or replace big, heavy boxes.

Think of it like the difference between hardware and software. Hardware is fixed — changing it requires intervention and downtime. Software is (or really *should be*) malleable, able to be easily reconfigured, updated, or extended. The old martech architecture behaved like hardware, even though it ran in the cloud. The new architecture delivers on the real promise of software: fluid, adaptable, continuously evolvable.

Standardize to diversify. This sounds paradoxical, but it's the key insight. By standardizing the foundation, you increase your ability to diversify what runs on top. A unified data platform doesn't constrain your choice of applications; it liberates it. When everything can access the same data through the same interfaces, you can mix and match capabilities freely. The explosion of apps and agents becomes manageable precisely because the underlying platform provides consistency.

Adjacency and adaptability as the payoff. Two benefits tower above others. First, *adjacency*: in a composable canvas, everything is closer to everything else. Data, services, apps, and agents are all accessible to each other without the friction of crossing application boundaries or building complex integration pipelines. Second, *adaptability*: when capabilities compose on a shared substrate rather than existing as solid boxes, change becomes fast, easy, and inexpensive. You can swap a service, add an agent, or reshape a workflow without heroic replatforming efforts.

5 Rings of Capability, One Center of Gravity

If the composable canvas isn't a vertical stack, does that mean it has no structure at all?

Not quite. The canvas is fluid in how capabilities connect and compose, but there's still an underlying architecture: a set of concentric responsibilities that radiate outward from a centralized data core to distributed apps and agents at the edge.

Think of it as five rings, each with a distinct role:

Data Core is the foundation: the unified repository of customer, company, content, code, and control data discussed earlier. This is the gravitational center of the architecture and the shared substrate on which everything else operates.

Semantic Layer sits above the core, although often bundled with it, standardizing definitions and governing usage. It's what makes the data meaningful and consistent across everything that touches it.

CaaS (Context-as-a-Service) platforms package relevant data and services for consumption by the outer rings. This is where CDPs, composable or otherwise, and traditional systems of record such as CRMs operate as engines that assemble and serve fit-for-purpose data and services to the apps and agents that need it. CaaS platforms take the arduous “engineering” work out of *context engineering*, accelerating value for all the apps and agents that use them.

Decisioning is where AI decisioning engines and reinforcement learning models optimize next-best-actions, and where orchestration controls resolve contention when multiple apps and agents all want to reach the same customer at the same time. Decisioning may be embedded into apps and agents, packaged with a CaaS platform such as a CRM or CDP, or run as a standalone service. These services prevent the outer ring from becoming a free-for-all.

Apps & Agents form the outermost ring, where customer experiences get delivered, insights surface for decision-makers, employees get work done, and business operations hum along through tailored workflows and automations. Commercially-packaged products still operate here, but increasingly this is hypertail territory: custom-built apps and agents where differentiation lives.

- **Apps & Agents** for customer and employee experiences, business operations; can be commercial products, but many will be custom-built
- **Decisioning** may be bundled in apps/agents, bundled with a CaaS platform, or run as a stand-alone service
- **CaaS** (context-as-a-service) platforms package relevant data and services to be used by apps/agents (e.g., CDPs, CRMs, and other coordinating systems of record)
- **Semantic Layer** standardizes a well-defined data catalog and governs proper usage
- **Data Core** is the repository of all data: customer, company, content, code, and control

*apps, agents, decisioning services, and CaaS platforms (such as a CDP) may be **embedded** natively on the data platform or **integrated** as a composed solution*

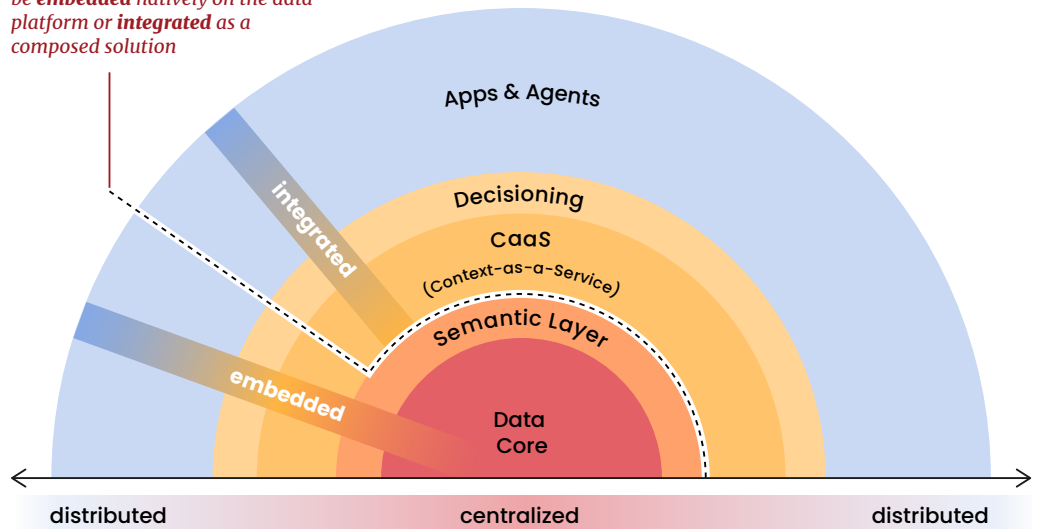


Figure 7.

Five rings of capability in the composable canvas

“I don’t want to call it a CDP anymore because it’s not necessarily a customer data platform. It includes everything from 360 to segmentation to activation to measurement and more importantly to data sharing.”

— Kumar Ram, VP/Global Head of Marketing Data Sciences, HP

The architecture shifts from centralized to distributed as you move outward. The inner rings of the Data Core and Semantic Layer benefit from centralization; that’s where consistency and governance matter most. The outer ring of Apps & Agents is inherently distributed, operating wherever customers and employees need them. CaaS platforms and Decisioning services fall somewhere in between: they can be more or less centralized depending on your architecture, but tend toward centralization within a given domain, such as marketing or sales.

Modern data platforms, such as Databricks, can host embedded capabilities across multiple rings natively: CaaS platforms (imagine a fully data-platform-native CDP), decisioning services, and apps and agents. But the composable canvas also accommodates these capabilities running elsewhere, seamlessly integrated with the data platform through the protocols and adjacency patterns we’ve discussed.

The structure is flexible; what matters is that everything connects.

There’s another way to read these rings: as an outward flow from *data* to *decisions* to *delivery*. The inner rings are fundamentally about data: storing it, governing it, assembling it. The middle rings are where decisions get made, informed by that data. And the outer ring of apps and agents is where delivery happens: experiences reach customers, tools enable employees, workflows execute operations.

Breaking the Integration Bottleneck

Integration has been the #1 technical challenge of martech for as long as martech has existed.

Ask any marketing ops professional about their biggest pain points, and integration will be near the top. Getting systems to talk to each other. Keeping data synchronized. Dealing with broken connections after vendor updates. Building custom pipelines because the out-of-the-box integrations don’t do what you need. The amount of time, money, and energy that has been poured into integration work across the industry is staggering.

But here’s the thing: how hard integration is depends on how systems are designed to relate to each other. And that design has been steadily evolving, with each architectural shift lowering the barrier a bit more.

In the 1st Age, integrations were predominantly *point-to-point*. System A needed to talk to System B, so you built a connection. System A also needed System C, another

¹⁰ Big O notation, Wikipedia — approximates the cost (in time/resources) for algorithms as n grows

connection. With n systems, you potentially needed $n(n-1)/2$ integrations. That's **$O(n^2)$ complexity**.¹⁰ Ten systems might mean 45 integrations. A hundred systems could mean nearly 5,000. This was clearly impractical, and as a result most systems remained unintegrated.

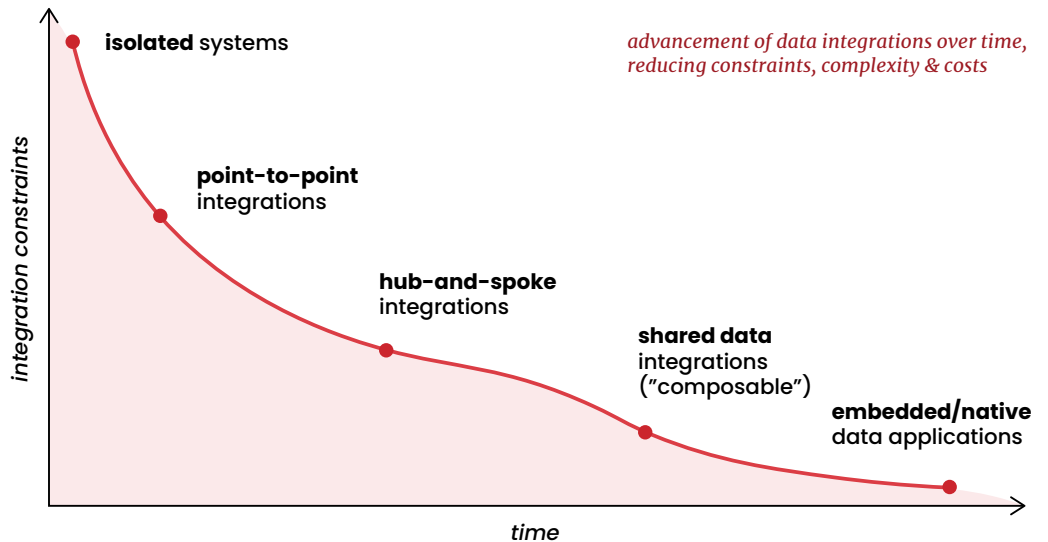


Figure 8.
The integration maturity curve: from silos to seamlessness

In the 2nd Age, platform ecosystems and *hub-and-spoke* architectures improved things. CDPs and major martech platforms with app marketplaces created central hubs. Each system only needed to connect to the hub, not to every other system. With n systems, you needed roughly n integrations — **$O(n)$ complexity**. A hundred systems meant a hundred integrations instead of thousands. Much better, but still a lot of overhead. The hubs themselves became bottlenecks: for throughput, for change management, for architectural flexibility.

In the 3rd Age, something different becomes possible. When data lives in a unified foundation and composable or embedded applications operate directly on that *shared data substrate*, the integration problem starts to fade. Systems don't need to exchange data with each other because they're inherently referencing the same data. Standardized protocols handle common interaction patterns. Complexity approaches **$O(\log n)$** , where each new capability adds minimal incremental integration burden because it's joining a coherent ecosystem, not connecting to a web of independent systems.

This isn't theoretical. It's the architectural payoff of everything discussed so far: the unified data foundation, the semantic layer, composable services operating on shared data. They combine to make integration fundamentally simpler.

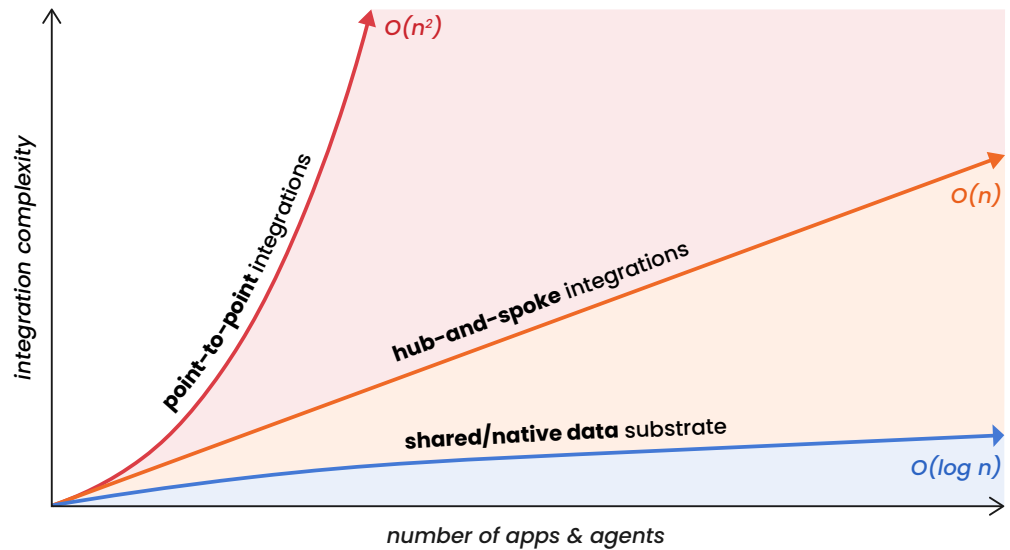


Figure 9.
Why integration architecture is a scaling decision, not just a technical one

The practical result: when a new channel emerges or a new tactic shows promise, you can act on it quickly. The martech stack stops being a constraint on marketing agility and starts being an enabler of it.

Composing Dynamic Partnerships

This architecture doesn't just simplify internal operations. It makes your external business ecosystem easier to activate.

Every company has an ecosystem: suppliers, distributors, agencies, technology vendors, channel partners, co-marketing allies, data providers. Historically, activating these relationships required significant technical work: negotiating data formats, building secure exchange mechanisms, reconciling schemas, maintaining pipelines. Partnership decisions often hinged as much on integration difficulty as on strategic value.

The composable canvas shifts that balance, letting commercial attractiveness drive decisions with fewer technical constraints.

¹¹ [Delta Sharing: An open standard for secure data sharing](#), The Linux Foundation Projects

For data sharing, open protocols like Delta Sharing¹¹ enable secure, real-time data exchange across organizational boundaries, including across different cloud infrastructure providers. A brand can share customer insights with a retail partner to optimize joint promotions. A publisher can provide audience data to advertisers for

better targeting. A manufacturer can give distributors real-time inventory visibility. Data clean rooms extend this for privacy-sensitive collaborations, enabling joint analysis without exposing underlying records. In all cases, the data stays where it is. No copying, no ETL pipelines, no reconciliation headaches.

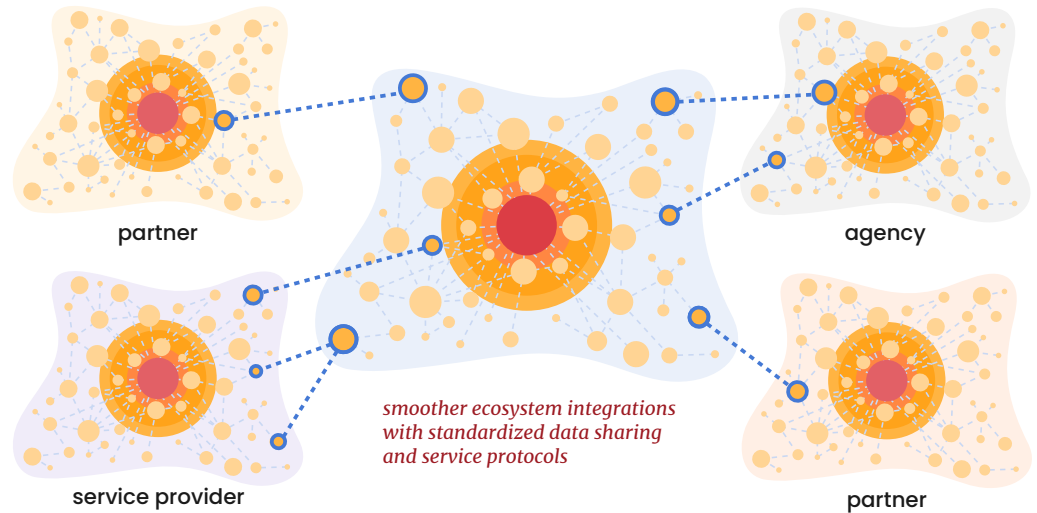


Figure 10.
Ecosystems and partnerships become composable

¹² *Under the Hood: Universal Commerce Protocol (UCP)*, Google, January 2026

For agent and service interoperability, the protocols discussed earlier — MCP, A2A, and emerging domain-specific standards like UCP for commerce¹² — collapse the timeline for activating partnerships. What used to take months of integration work can happen in days or hours. A brand could spin up a co-marketing campaign with a new retail partner: agents coordinating on audience targeting through A2A, enabling direct checkout with promotional pricing through UCP, processing payments through whichever handler the customer prefers. When the promotion ends, the partnership can wind down cleanly. No orphaned integrations. No lingering pipelines to maintain.

These capabilities ripple across partnership types:

- **Agencies and service providers** can plug directly into client workflows rather than exchanging files and sitting in status meetings.
- **Channel and go-to-market partners** can access approved content, product information, and customer insights through standardized interfaces.
- **Data and technology partners** become more composable — integration effort drops from a major factor in partner selection to a minor consideration.

Ecosystems and partnerships become composable: new combinations emerge without heavy integration barriers or bottlenecks.

“We want to own the core construct of the data and infrastructure. That’s why having it in Databricks is the right choice. If we change agencies, all we need to do is flip the activation layer at the top. The foundation stays with us.”

— Kumar Ram, VP/Global Head of Marketing Data Sciences at HP

Getting There From Here

The composable canvas is not a product you purchase or a project you complete. It's an architectural direction, a way of thinking about how marketing technology should evolve.

The good news: you don't need to tear down your existing stack. The journey is incremental, and every step can deliver value on its own terms. Different organizations will take different paths depending on their starting point, their capabilities, and their priorities. But certain principles apply broadly.

Start With the Data Foundation

The single highest-leverage investment is getting your data house in order.

If you don't already have a modern data platform, that's where to begin. This isn't a marketing decision. It's an enterprise infrastructure decision that marketing should be advocating for loudly. Marketing use cases often justify the investment, but the platform serves the entire organization.

“Granular data access control lets us treat the whole system as one ecosystem with different contextual views into it — rather than duplicating data into separate systems.”

— Chris Wissing, Chief Product Officer, Epsilon

If you already have a data platform, the question becomes: is your marketing data actually in it? For many organizations, the honest answer is “some of it.” Campaign performance lives in your marketing automation platform. Customer interactions are scattered across CRM, support system, and product analytics. Website behavior sits in a separate analytics tool. Ad performance reports are downloaded as spreadsheets.

The first step is data consolidation: getting marketing data flowing into the platform where it can join with everything else. In your current stack, this is likely traditional plumbing work: connectors, pipelines, ETL jobs, data quality checks. Not exciting, but foundational. And it delivers immediate value: unified reporting, cross-system analysis, a single version of truth. More importantly, it ends the reporting wars. When everyone — marketing, sales, finance, the board — is looking at the same numbers, calculated the same way, meetings shift from debating data to deciding what to do about it.

Build the semantic layer. As data consolidates, invest in meaning. A semantic layer defines what your data actually represents in business terms. What constitutes a

“We’ve designed our stack like modular Lego blocks so the business can move faster. As new technologies and AI agents emerge, we have the flexibility to adopt the best tools with minimal lift from a lean team. That portability gives us real agility — we can evolve the stack as the business evolves, without heavy migrations or brittle integrations.”

— Elizabeth Dobbs, AVP of Marketing Technology, Data & Growth, Databricks

“customer?” How do you calculate “lifetime value?” What does “engaged” mean for your business?

Start with your most critical concepts: customer definition, key events (purchase, signup, churn), core metrics (revenue, conversion, engagement). Document them. Implement them in your data platform so queries reference canonical definitions rather than ad-hoc logic. This is harder than it sounds. It requires cross-functional agreement on things people thought they agreed on but didn’t.

Embed governance from the start. As you build the foundation, embed governance. Who can access what data? What are the privacy constraints? How do you track lineage when data flows through transformations?

Modern data platforms provide excellent governance infrastructure — Databricks’ Unity Catalog is one example. Use it. Tag sensitive data. Implement access controls. Establish audit trails. This isn’t just compliance overhead. It’s what allows you to move faster later, because you can confidently extend access to new tools and teams knowing guardrails are in place.

Composability as Your North Star

Every technology decision you make — every new tool you evaluate, every integration you build, every vendor you select — should be filtered through a simple question: *Does this move us toward composability or away from it?*

Favor openness. When choosing between vendors, prefer those who embrace open standards, open data formats, and open APIs. A platform built on proprietary data formats creates lock-in; one built on Delta Lake or Iceberg lets you change your mind later. A tool that exposes capabilities through well-documented APIs can participate in a composed architecture; one that guards its functionality behind a UI cannot.

This doesn’t mean rejecting proprietary software. Many excellent tools are proprietary. It means understanding where proprietary choices create dependencies and making those choices deliberately rather than by default.

Prefer adjacency. When evaluating new capabilities, ask: can this operate close to our data, or does it require moving data elsewhere? Native data platform solutions that query your data directly align with the composable model. Solutions that require syncing data into their own proprietary store add complexity, introduce latency, and create governance challenges.

This principle applies particularly to customer data platforms. Traditional CDPs built their value proposition on data consolidation, bringing data from many sources into their platform. Composable CDPs flip this: they bring CDP capabilities (identity resolution, audience building, activation) to data that stays in your data platform. The composable approach leverages your existing data infrastructure investment rather than duplicating it.

Plan for replaceability. Assume that any tool you implement today might be replaced in three years. That's not cynicism. It's realism about how fast this market moves. Design integrations that minimize switching costs. Document your configurations. Avoid "overfit" customizations that couple your processes to a specific vendor's implementation details.

The goal is optionality. A composable architecture maximizes your ability to swap components, try new tools, and evolve your stack without massive migration projects.

From Boxes to Building Blocks

The composable canvas doesn't eliminate existing martech categories, but it reframes their role. Understanding where your current tools fit helps you get more from what you've already built and add new capabilities without starting over.

Data foundations: platforms that consolidate and govern. Your data platform (Databricks, Snowflake, BigQuery, etc.) becomes the foundation. Customer data platforms (CDPs) serve as CaaS (context-as-a-service) and decisioning platforms in our 5 Rings of Capability, providing identity resolution, audience segmentation, predictive scoring, and other context-engineering capabilities for marketing, cleanly packaged for your campaigns and channels. Master data management (MDM) and product information management (PIM) systems feed clean, governed data into the foundation and keep it current.

CDPs vary in how tightly they connect to your data platform. Composable CDPs operate directly on your data platform, with the benefits of no duplication and everything governed in one place. But all CDPs deliver value by productizing context-engineering capabilities that would otherwise require significant custom development.

Intelligence and decisioning: platforms that analyze and decide. Business intelligence tools enable humans to explore data and surface insights. Natural language interfaces now let users "talk to their data," asking questions in plain

English rather than writing queries or navigating dashboards. This is also where next-best-action engines, propensity models, and AI decisioning platforms operate.

Established platforms like Power BI and Tableau continue to evolve with AI-powered capabilities, while modern data platforms such as Databricks AI/BI increasingly offer embedded analytics natively. The big shift in this category is BI's relationship to action. In the old model, analysts produced reports that humans interpreted and translated into campaign decisions. In the composable model, insights increasingly flow directly into decisioning systems, and AI agents can act on those insights autonomously, not just surface them. The analyst's role shifts from producing reports to designing the metrics, models, and feedback loops that intelligent systems use.

Customer engagement: platforms that execute and deliver. Marketing automation platforms (MAPs), email service providers (ESPs), and customer engagement platforms (CEPs) remain essential. They're still how you reach customers across channels. These platforms have invested heavily in delivery infrastructure, channel management, compliance capabilities, and campaign development tools that you shouldn't replicate.

What's evolving is how these platforms participate in a broader ecosystem. Traditionally, the MAP or CEP operated as a self-contained environment: data came in, marketers built segments and journeys within the platform, messages went out. In the composable model, these platforms become more interoperable, able to receive audiences, triggers, and decisioning signals from upstream systems while contributing their own intelligence and capabilities.

The vendors themselves are building deeper integrations with data platforms, exposing more functionality through APIs, and enabling more flexible architectures. The question for your organization is how to take advantage of that flexibility: which decisions happen where, and how do the pieces connect?

Experience delivery: platforms that render and present. Content management systems (CMS) and digital experience platforms (DXP) manage what customers see. Digital asset management (DAM) systems store and organize creative content. These platforms shape the moments that matter most: actual interactions customers have with your brand. And they're becoming more composable.

First, headless architectures that separate content management from delivery — a headless CMS delivers content through APIs to any front-end, enabling the same content to render across web, mobile, in-app, and email without duplication.

Second, content itself increasingly becomes a governed asset within the data platform, carrying metadata describing what it is, who it's for, and how it performs.

Media and advertising: platforms that buy and measure. Demand-side platforms (DSP) and ad servers handle programmatic advertising. Attribution platforms measure what's working. These tools connect to walled gardens (Google, Meta, Amazon) that aren't going away.

The composable model changes how data flows to and from these systems. Clean rooms and privacy-safe data sharing protocols enable richer audience activation and measurement without exposing sensitive customer data. They also facilitate enrichment through third-party identity vendors and connections to retail media networks. First-party data from your unified foundation becomes the basis for ad targeting, suppression, and closed-loop measurement, reducing wasted spend and proving impact.

“I do think SaaS solutions are in trouble if they aren't bringing meaningful leaps forward for marketing teams. If you're just accessing data in an easy way, we're going to do that ourselves. So you've really got to build something defensible on top at the experience layer.”

— Meagen Eisenberg,
CMO, Samsara

Integration and automation: platforms that connect and orchestrate. These tools traditionally spent much of their effort replicating data between systems, keeping customer records in sync across CRM, email, and ad platforms. In the composable canvas, where data stays in the unified foundation, their role evolves from moving data to routing signals and governing orchestration.

Integration platforms become the nervous system of your martech environment. A high-value customer submits a support ticket, and a signal cascade follows: pause promotional emails, alert the account team, suppress retargeting ads. The customer data stays where it is; only the signals flow to the systems that need to act.

Workflow automation evolves from rigid *if-then* rules to intelligent orchestration. AI agents monitor campaign performance, reallocate budget when channels underperform, and adapt sequences based on how customers respond. Critically, these platforms can also provide governance, defining which agents can take which actions, what approvals are required, and how decisions are logged for audit and continuous improvement.

Existing platforms and new capabilities will coexist for years. That's not a compromise — it's prudent. Manage the hybrid state intentionally: know which capabilities run where, let new approaches prove themselves, and migrate deliberately.

The goal is controlled evolution: incremental progress toward the North Star, with each step delivering value and reducing risk for the next.

Clarity, not chaos. Progress, not disruption.

Conclusion

The martech stack served us well for two decades. It gave structure to a chaotic proliferation of tools, provided a mental model for thinking about how systems relate, and guided billions of dollars in technology investment.

But the stack metaphor is showing its age. The rigid layers, the vertical structure, the headaches of integration. These reflect an era when martech was about buying boxes and plumbing them with pipes.

“In the Mad Men era, you started with the big idea. Now you can flip that. Start with the data to identify your most important personas, then let the creative flow from there. The big idea doesn’t disappear, but it’s informed differently.”

— Chris Wissing, Chief Product Officer, Epsilon

The 3rd Age of Martech invites a different metaphor: the composable canvas. It’s an architecture where data is the foundation everything else builds on, where capabilities can be assembled and reassembled fluidly, where agents and applications operate on a shared substrate that makes integration nearly invisible.

This isn’t a utopian vision. It’s a practical direction, grounded in technical capabilities that exist today and will mature significantly over the next 3–5 years. Modern data platforms, open protocols, AI-powered orchestration, and the ease of building custom software make this future achievable for organizations willing to invest in it.

The journey will be different for every company — different starting points, different priorities, different constraints. There is no single path. But that’s the beauty of composability. It bends to your business, not the other way around.

The direction is clear: the organizations that treat composability as a North Star, invest in their data foundation, and build teams capable of operating in this new environment will have turned their martech from a constraint into a competitive advantage.

The 3rd Age is beginning. Time to start painting your canvas.

About Databricks

Databricks is the Data and AI company. More than 20,000 organizations worldwide — including Block, Comcast, Condé Nast, Rivian, Shell and over 60% of the Fortune 500 — rely on the Databricks Data Intelligence Platform to take control of their data and put it to work with AI. Databricks is headquartered in San Francisco, with offices around the globe and was founded by the original creators of Lakehouse, Apache Spark™, Delta Lake, MLflow, and Unity Catalog. To learn more, follow Databricks on [X](#), [LinkedIn](#) and [Facebook](#).

About the Author

Scott Brinker, dubbed the “Godfather of Martech” by *Ad Age*, has been analyzing marketing technology and its impact on marketing organizations for more than 18 years as editor of [chiefmartec](#). He previously served as VP of Platform Ecosystem at HubSpot and before that was co-founder and CTO of ion interactive, a pioneering no-code platform for interactive content. Scott authored the best-selling book *Hacking Marketing* and co-authored *The New Automation Mindset*. He holds graduate degrees in computer science and management from Harvard and MIT. For his latest analysis and insights on martech, subscribe to his [newsletter](#) and follow him on [LinkedIn](#).