PART I: THE GROUND SHIFTED

CHAPTER ONE

The Three Eras of Enterprise Data

There is a company in Stuttgart—we can't name them, but you'd recognize the logo—that has been running SAP for thirty-one years. Their SAP Business Warehouse holds over two decades of transactional history: every purchase order, every invoice, every production run, every quality inspection, every customer interaction, organized into cubes and InfoProviders that a generation of consultants built and a generation of analysts learned to query. When we sat down with their CIO in late 2024, she said something that has stayed with us ever since: "We have more data than any company in our industry. And we have less insight than a startup with a Snowflake account and two data engineers."

That sentence captures the paradox that has defined enterprise data management for the last decade, and it's the reason this chapter exists. The history of how companies collect, organize, and activate data is not just a technology story—it's the story of what becomes possible when the architecture changes. Every time the data foundation shifts, the go-to-market motion shifts with it. The companies that recognize the shift early build advantages that compound for years. The ones that wait get commoditized—not because they lack talent, but because they're running a modern playbook on yesterday's infrastructure.

We think of this history in three eras, defined not by the marketing tools that sat on top but by the data architecture that sat underneath everything. Not the campaigns. Not the channels. Not the org chart. The architecture—because the architecture determines what's possible, and what's possible determines how you go to market.

The history of how companies collect, organize, and activate data is not just a technology story — it's the story of what becomes possible when the architecture changes. Every time the data foundation shifts, the go-to-market motion shifts with it.

> *"We have more data than any company in our industry. And we have less insight than a startup with a Snowflake account and two data engineers."*

Figure 1: The Three Eras of Enterprise Data

Three eras, defined not by the marketing tools on top but by the data architecture underneath. The architecture determines what's possible, and what's possible determines how you go to market.

Dimension	Era 1: Warehouse 1990s–2010s	Era 2: Data Lake 2010s–2020s	Era 3: Data Fabric 2020s–present
The tagline	***"A DMP knew your audience."***	***"A CDP knows your customer."***	***"A data fabric knows your business."***
Architecture	Data warehouse (SAP BW, Teradata, Oracle). ETL: move data to analytics. Batch overnight.	Data lakes, lakehouses, CDPs. Fragmented — CRM, ERP, supply chain in separate systems. Unification problem.	Business data fabric (SAP BDC). Bring intelligence to the data. Semantic layer preserves context. Zero-copy exchange.
Data paradigm	Structured. Predefined questions. Periodic reports — monthly, quarterly, annually.	Unified customer view. Real-time signals. Identity resolution across systems.	Federated. Contextual. Agents reason across operational data with semantic understanding intact.
Who uses it	Small priesthood of analysts and finance. Everyone else gets dashboards or slide decks.	Broader access. Marketers, product teams, CX. Self-serve analytics begins.	AI agents as primary consumers. Humans direct; agents execute and synthesize.
PMM job	Static deliverables: positioning docs, messaging frameworks, campaign briefs. Quarterly cadence.	Expanded: lifecycle marketing, journey mapping, win/loss, persona development, data-driven pricing.	Redefined: from executor who produces artifacts to strategist who directs agents producing artifacts.
PMM skill set	Communications-heavy. Storytelling + enough technical fluency to be credible.	Data fluency. Credible with a CTO on data integration, explain it to a CMO in business terms.	Systems orchestration. Which parts automate, which need judgment, how to design the handoff.

> **The shift:** The earlier transitions asked PMMs to expand their skill set. The fabric-to-agent shift asks PMMs to redefine their role entirely — from executors who produce artifacts to strategists who direct agents that produce at a pace no individual human could match.

Figure 1. The Three Eras of Enterprise Data. Each era reshaped not just technology stacks but organizational talent models. The companies that recognize the shift early build advantages that compound for years.

We think of this history in three eras, defined not by the marketing tools that sat on top but by the data architecture that sat underneath everything.

Figure 1: The Three Eras of Enterprise Data

Era One: The Data Warehouse and the Age of Reports

For most of the 1990s and 2000s, the center of gravity for enterprise data was the data warehouse. SAP BW, Teradata, Oracle—these were the

systems of record for analytics, and they operated on a simple premise: extract data from your operational systems, transform it into a structured model, and load it into a warehouse where analysts could run reports. ETL—extract, transform, load—was not just a technical process; it was a philosophy. You moved the data to the analytics. You waited for the batch job to run overnight. And in the morning, you had your numbers.

The data warehouse era was, in hindsight, wonderfully contained. The data was structured. The questions were predefined. The reports were periodic—monthly, quarterly, annually. And the people who consumed the data were a small priesthood of analysts and finance teams who knew how to write queries and build Excel models. The rest of the organization got dashboards, if they were lucky, or slide decks summarizing what the dashboards said.

For product marketers, this era defined the job in ways that persist to this day. The PMM's deliverables were fundamentally static: positioning documents, messaging frameworks, campaign briefs, competitive battlecards that were accurate on the day they were written and started decaying the moment they were saved. The Pragmatic Institute built its framework during this era, and the framework reflects the cadence of warehouse-era data—quarterly reviews, annual planning cycles, periodic competitive updates. The data moved slowly, so the marketing moved slowly, and that was fine because the buyers moved slowly too.

Chris O'Hara lived the marketing side of this era at Krux, one of the pioneering data management platforms that Salesforce later acquired. Even in the adjacent world of advertising data, the same warehouse-era assumptions held: you collected audience data in batches, segmented it into taxonomies, and activated it against campaigns that ran for weeks or months. The DMP—data management platform—was the warehouse paradigm applied to marketing, and it worked until it didn't.

Era Two: The Data Lake and the Age of Unification

The second era arrived not as a single product launch but as a slow-motion crisis of fragmentation. Enterprises woke up one morning—metaphorically, though some of them describe it exactly that way—and realized their data was everywhere. The warehouse still existed, but now there was also a data lake on AWS. A Snowflake instance that the analytics team had spun up. A Databricks workspace for the data science team. Half a dozen SaaS applications, each with its own data model and its own API. CRM over here, ERP over there, supply chain data in a third system, marketing data in a fourth. The dream of a single source of truth had shattered into a hundred sources of partial truth, none of them quite agreeing with the others.

This was the era that produced the customer data platform—the CDP—which Chris and Martin Kihn wrote the book on (literally: *Customer Data Platforms*, Wiley, 2021). The CDP was an attempt to solve the unification

problem for marketing: take all the customer data from all the systems, resolve the identities, and create a single customer view that marketers could actually use. The tagline that emerged from that work captured the generational shift: "A DMP knew your audience. A CDP knows your customer."

But the unification problem was never just a marketing problem. It was an enterprise problem. Every function—finance, supply chain, HR, procurement—was drowning in the same fragmentation. And the solutions that emerged reflected this: data lakehouses that tried to combine the structure of warehouses with the flexibility of lakes, data mesh architectures that tried to decentralize ownership while maintaining governance, and—most importantly for our story—the concept of a business data fabric that could connect data across systems without requiring you to move it all to one place.

This is where SAP's journey becomes central to the story, because SAP sits at the intersection of the problem in a way no other company does. SAP runs the operational backbone of the world's largest enterprises—the ERP, the supply chain, the financials, the HR systems. When we launched SAP Datasphere in 2023, the product strategy was built around a recognition that had been forming for years: customers didn't need another place to put their data. They needed a way to make their data meaningful wherever it already lived. Datasphere introduced the business data fabric—a semantic layer that preserved the business context of SAP data and connected it to non-SAP sources without forcing costly extraction and replication. It was, in retrospect, the bridge between the unification era and what came next.

For product marketers, the unification era expanded the job dramatically. The PMM was no longer just writing positioning and campaign briefs. Now the job included lifecycle marketing, personalization strategy, customer journey mapping, win/loss analysis, and a whole set of activities that required understanding individual customers and their relationship with the product. The Pragmatic Institute's framework grew to accommodate this—more activities around customer insights, persona development, data-driven pricing. The PMM became less of a campaign strategist and more of a customer strategist.

Era Three: The Data Fabric and the Age of Agents

Which brings us to now. And the line that captures this era: a DMP knew your audience. A CDP knew your customer. A data fabric—powered by agentic AI—knows your business.

In February 2025, we launched SAP Business Data Cloud—a fully managed SaaS platform that unified SAP Datasphere, SAP Analytics Cloud, and SAP Business Warehouse with an embedded Databricks partnership, all built around a shared object store and zero-copy data exchange. The product strategy behind BDC was something we had been building toward

for years: stop asking customers to move their data to the analytics, and instead bring the intelligence to the data. Build semantic data products that carry their business context with them—so that "Customer" in SuccessFactors and "Customer" in S/4HANA don't get flattened into the same field by a junior data engineer at 2 AM. Federate the metadata. Keep the data where it lives. And make the whole thing the foundation for AI that actually understands the business it's operating in.

We need to be careful here, because "agentic" is rapidly becoming one of those words that means everything and nothing, like "digital transformation" or "cloud-native." When we say agentic, we mean something specific: software systems that can perceive their environment, make decisions, and take actions without waiting for a human to tell them what to do at each step. Not chatbots. Not copilots that suggest the next sentence in your email. Agents that monitor your competitive landscape while you sleep. Agents that adjust your pricing recommendations based on real-time market signals. Agents that generate and test messaging variants based on what's resonating with different buyer segments. Agents that evaluate your demo environment and flag when it's showing a workflow that doesn't match the customer's use case.

The data fabric sits at the convergence of three things that didn't exist—or didn't exist at sufficient scale—five years ago. First, large language models that can reason about unstructured business context in ways that previous AI could not. Second, rich enterprise data fabrics—like SAP Business Data Cloud—that give those models access to the operational data (ERP, CRM, supply chain, finance) that actually runs the business, with the semantic context intact. Third, agent orchestration frameworks that allow multiple AI agents to collaborate on complex, multi-step workflows without constant human supervision. HANA Cloud's knowledge graph, its multi-modal compute engine spanning relational, document, graph, and knowledge graph processing—these are not features on a product roadmap. They are the infrastructure that makes enterprise-grade agentic AI possible.

The data fabric sits at the convergence of three things that didn't exist — or didn't exist at sufficient scale — five years ago. Together, they make enterprise-grade agentic AI possible.

Figure 2: The Agentic Convergence

Three capabilities converging to create the agentic era. None is sufficient alone. All three together change who does the work.

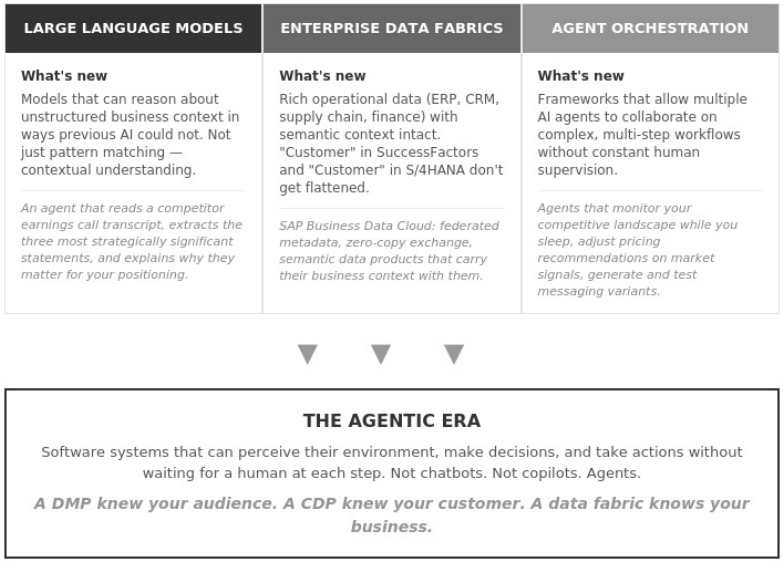

Figure 2. The Agentic Convergence. LLMs provide reasoning. Data fabrics provide context. Agent orchestration provides autonomy. Together, they change not what data you have access to, but who does the work.

"Agentic" is rapidly becoming one of those words that means everything and nothing. When we say agentic, we mean something specific: software systems that can perceive, decide, and act without waiting for a human at each step.

Figure 2: The Agentic Convergence

For product marketers, this changes the game in a way that's different from the warehouse-to-lake transition. That transition changed what data you had access to and what you could do with it. This transition changes who does the work. The earlier shifts asked PMMs to expand their skill set. The fabric-to-agent shift asks PMMs to redefine their role entirely—from executors who produce artifacts to strategists who direct agents that produce artifacts at a pace and quality that no individual human could match.

That is an uncomfortable sentence to write, and probably an uncomfortable one to read. But we've watched enough hype cycles to know that the uncomfortable truths are the ones worth paying attention to early.

Each data era reshaped not just the technology stack but the PMM's role, deliverables, and required skill set. The progression is clear — and the third shift is the most disruptive.

Figure 3: The PMM Role Shift Across Eras

How the PMM job changed with each architectural shift. The first two transitions expanded the job. The third redefines it.

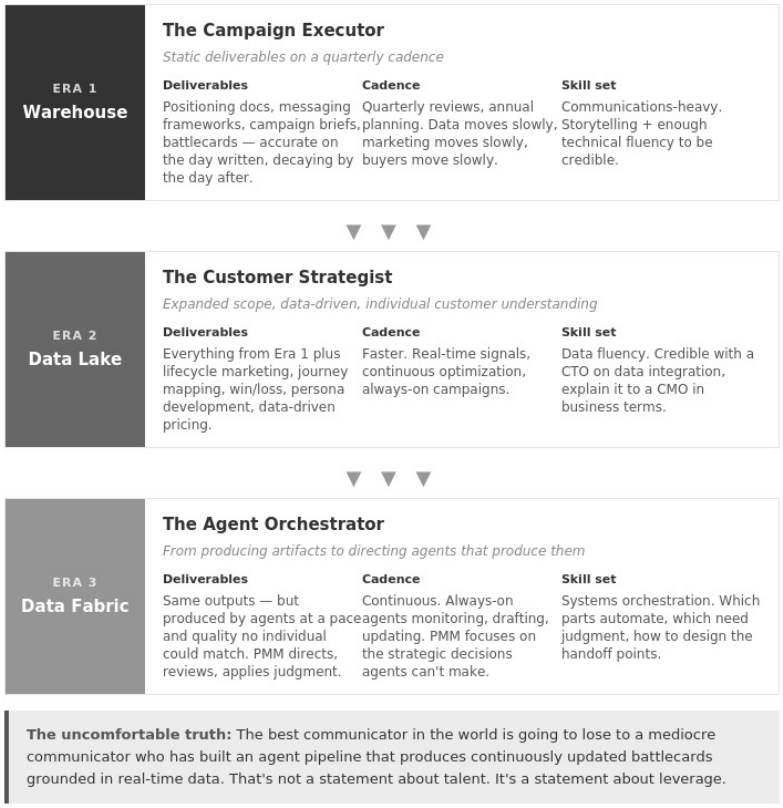

Figure 3. The PMM Role Shift Across Eras. Era 1 expanded to Era 2. Era 2 expanded to Era 3. But Era 3 is different — it doesn't just add activities. It changes who does them.

The fabric-to-agent shift asks PMMs to redefine their role entirely — from executors who produce artifacts to strategists who direct agents that produce artifacts at a pace and quality that no individual human could match.

Figure 3: The PMM Role Shift Across Eras

• • •

The rest of this book is organized around a simple question: for each major activity in the product marketing job, what changes when agents become part of the team? Not what changes in theory. What's already changing, right now, for the PMMs who are paying attention.

But before we get there, we need to do something that we think nobody has done honestly yet: take the canonical framework for product marketing and hold it up to the light. Chapter 2 is about the Pragmatic Institute's thirty-seven activities, and what happens when you score each one for AI disruption potential. The results are sobering—and, if you're inclined

toward optimism, genuinely exciting.

The CMO's perspective maps each era to an organizational architecture and hiring profile. The progression has moved from communications to data fluency to systems orchestration — and leverage, not effort, is the new differentiator.

Figure 4: The Leverage Equation

What CMOs need from PMMs in each era — and why the agentic era demands a fundamentally different talent model.

Dimension	Era 1: Warehouse	Era 2: Data Lake	Era 3: Data Fabric
Hiring profile	Communications-heavy. Storytelling plus enough technical fluency to be credible.	Data fluency. Credible with a CTO on integration, explain it to a CMO in business terms.	Systems thinker. Can look at a complex workflow and figure out what to automate vs. what needs judgment.
Differentiator	Quality of writing and messaging. The best communicator wins.	Depth of customer understanding. The most data-fluent PMM wins.	Leverage. The PMM who builds agent pipelines beats the one who manually produces artifacts.
Org design	Well-defined roles: messaging, positioning, campaign support. Functional silos.	Broader roles. PMMs needed across lifecycle, personas, pricing, analytics. Cross-functional.	Orchestration roles. PMMs design workflows, direct agents, apply judgment at handoff points.
Capacity model	Linear. Output = headcount × hours. More people = more output.	Partially leveraged. Better data = better targeting = more efficient campaigns.	Exponential. One PMM with agent pipelines outproduces a team without them.
Risk of being late	Slow decay. Competitors with better messaging gradually win positioning.	Faster erosion. Data-driven competitors personalize while you're still in quarterly cycles.	Structural gap. Competitors with agent leverage produce 10x while you produce 1x. The gap compounds.

Leverage, not effort, is the new differentiator.
Building agent pipelines beats manually producing artifacts.

> **For the book ahead:** Every remaining chapter asks a version of the same question — for each major activity in the PMM job, what changes when agents become part of the team? Not in theory. What's already changing, right now, for the PMMs who are paying attention.

Figure 4. The Leverage Equation. Each era shifted the talent model. In the agentic era, the differentiator isn't the quality of what you produce — it's the leverage of how you produce it. The gap between agent-enabled PMMs and manual-execution PMMs compounds over time.

But before we get there, we need to do something nobody has done honestly yet: take the canonical framework for product marketing and hold it up to the light.

Figure 4: The Leverage Equation

THE CMO PERSPECTIVE

The three-era framework maps not just to data architecture but to organizational architecture. In the warehouse era, the PMM job was well-defined: messaging, positioning, campaign support. The skill set was

communications-heavy—storytelling plus enough technical fluency to be credible. In the unification era, the profile shifted. You needed people who could have a credible conversation with a CTO about data integration and identity resolution and then explain it to a CMO in business terms. The job got harder, the skill set got broader, but it was still fundamentally a human-execution job.

In the agentic era, the hiring profile changes again. What we need now are people who can orchestrate—who can look at a complex workflow like building a competitive battlecard and figure out which parts should be automated, which parts need human judgment, and how to design the handoff points. That's a systems thinking skill, not a communications skill. The best communicator in the world is going to lose to a mediocre communicator who has built an agent pipeline that produces continuously updated battlecards grounded in real-time data. That's not a statement about talent. It's a statement about leverage.

KEY TAKEAWAYS

- Each data era—warehouse, lake, fabric—reshaped not just technology stacks but organizational talent models.
- The PMM skill progression has moved from communications to data fluency to systems orchestration.
- Leverage, not effort, is the new differentiator: building agent pipelines beats manually producing artifacts.
- The agentic era demands that every PMM understand where human judgment ends and automation begins.