WHITE PAPER

The Case for Ontology as the Foundation of AI-driven Marketing

Infosys aster

Executive Summary

Enterprise marketing stacks are highly connected, yet AI tools produce strategically shallow outputs because they lack a shared understanding of what data means.

The missing layer is ontology: a structured map of relationships between customers, personas, journeys, content, and intent. While databases store facts, ontologies store meaning.

This paper outlines a five-stage process for building an ontology and integrating with Customer Data Platforms (CDPs), lead scoring, journey orchestration, and analytics via a Knowledge Graph. Ultimately, it argues that this contextual intelligence determines whether marketing AI delivers true business value.

Connected Stack, Disconnected Intelligence

McKinsey's Global Survey on the State of AI states that Gen AI is most used in marketing and sales across all industries surveyed, with 71% of companies using it in at least one business function.[1] Enterprise marketing teams in 2026 operate at a level of data coverage that would have seemed excessive a decade ago. The average enterprise marketing stack pulls in data from dozens of live sources simultaneously: Customer Relationship Management (CRM) records, web analytics, ad platform signals, email engagement data, social interactions, intent data providers, and third-party audience enrichment feeds. AI-powered tools sit on top of this infrastructure to score leads, manage journeys, personalize content, and optimize campaign spend.

On paper, the architecture looks like it should work and the volumes of captured data continue to grow quarter on quarter. But the questions that matter most to marketing practitioners remain difficult to answer.

- Why did this campaign convert one segment but not another?
- Why is the lead scoring model assigning high intent to prospects who never close?
- Why are AI recommendations inconsistent across channels when all the tools are reading from the same CRM?

Most marketing AI implementations are fully connected. The intelligence, however, is still shallow.

The Problem with AI That Doesn't Understand Context

Current AI systems deployed in marketing stacks are good at one thing: recognizing statistical patterns in historical data and projecting them forward. They can identify that a content type tends to produce higher click-through rates with a given audience segment. They cannot determine whether those clicks are advancing revenue or simply attracting the wrong audience more efficiently.

A lead scoring model operating on engagement frequency will assign high intent to any account that clicks a lot, regardless of whether those clicks come from a genuine buyer or an intern doing background research. A journey orchestration tool routing by historical conversion rates will send every high scorer to the same nurture sequence, irrespective of whether their buying stage or buying role warrants it. A personalization engine will surface content based on behavioral similarity, whether or not that content is strategically appropriate for where the account sits in the funnel.

Collectively, the stack is making decisions without a shared understanding of what anything means.

Every new AI tool added without a shared semantic foundation introduces another interpretation of the same data. 'Lead' means something in the CRM. It could mean something slightly different in the marketing automation platform. And something else again in the ad platform's audience model. AI does not resolve these definitional conflicts. It absorbs them, encodes them, and scales them.

No single output looks wrong, but the pattern shows up when you step back and observe it as a whole. Campaigns perform on vanity metrics but fail to move the pipeline; AI recommendations contradict each other across channels, and teams face a persistent inability to answer the questions that senior leadership actually asks.

Ontology - How AI Learns What Data Actually Means

The word "ontology" has its roots in philosophy. It is the study of what exists and how things relate to one another. In computer science, an ontology is a shared semantic resource model: a structured map of relationships between entities such as customers, accounts, personas, content, journeys, intent, and outcomes within a domain.

An ontology is a map of meaning. Where a database stores facts, such as clicks, revenue, and reach, an ontology stores the relational logic that gives those facts purpose. A customer belongs to an account. A persona influences a buying decision. An engagement signal indicates intent. A journey stage determines the relevance of a content type. These relationships form the shared truths that allow AI systems to reason across data silos, rather than merely calculating probabilities within them.

For instance, a Director of Demand Generation at a Series B company visiting a pricing page during a quarterly planning cycle, is in a different buying context than a graduate student clicking the same link. The AI system that understands that difference will make better decisions. This understanding comes from ontology.

The precedent for this is well-established at scale. Google's Knowledge Graph, an ontological map of over 500 billion facts about 5 billion entities,[2] improved search quality by adding relational meaning to raw data. Meta's social graph is a well-known example of how relationships between people, interests, content, and places can be represented and used to provide more relevant and personalized experiences.

Ontology in Digital Marketing

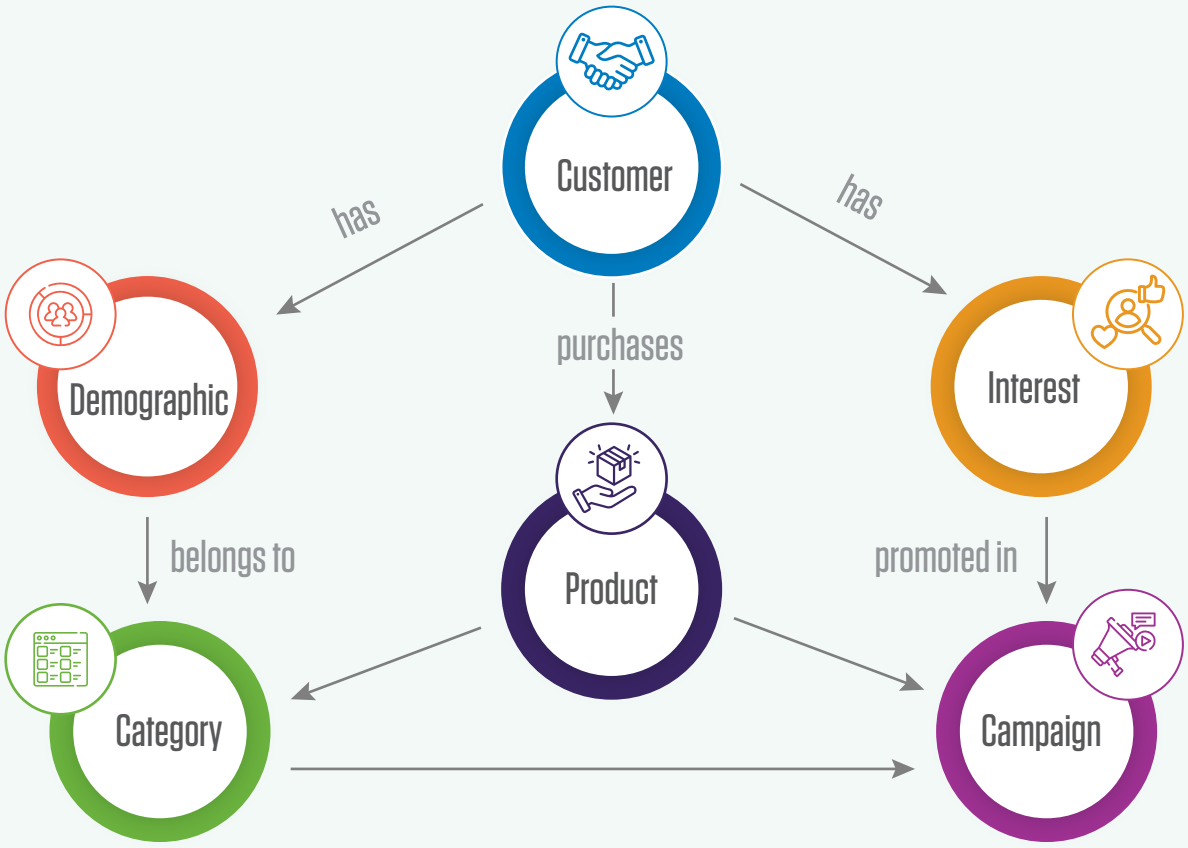

The Three-layer Architecture

AI-driven marketing will be built on three working layers that must work together for superior outcomes.

- **Layer 1 – Data**
  The raw inputs: individual and account-level signals, first-party behavioral data, second-party partner intelligence, third-party market data, and marketplace inputs. Most enterprises have this layer in reasonable shape.

- **Layer 2 – AI**
  The reasoning and prediction engines: models that generate synthetic data to fill gaps, reason and make predictions about customer behavior, and automate campaign decisions. The marketing industry has invested heavily in this layer over the past three years.

- **Layer 3 – Ontology**
  The meaning layer: the structured representation of what data actually signifies, how entities relate to one another, and why a given signal matters in a given context. Almost no enterprise marketing stack has this layer today.

Without the connective tissue that Layer 3 provides, Layers 1 and 2 are working at a structural disadvantage. AI with access to data but without shared meaning will produce sophisticated-sounding outputs that lack strategic coherence.

The 3-layer Architecture for AI-driven Marketing

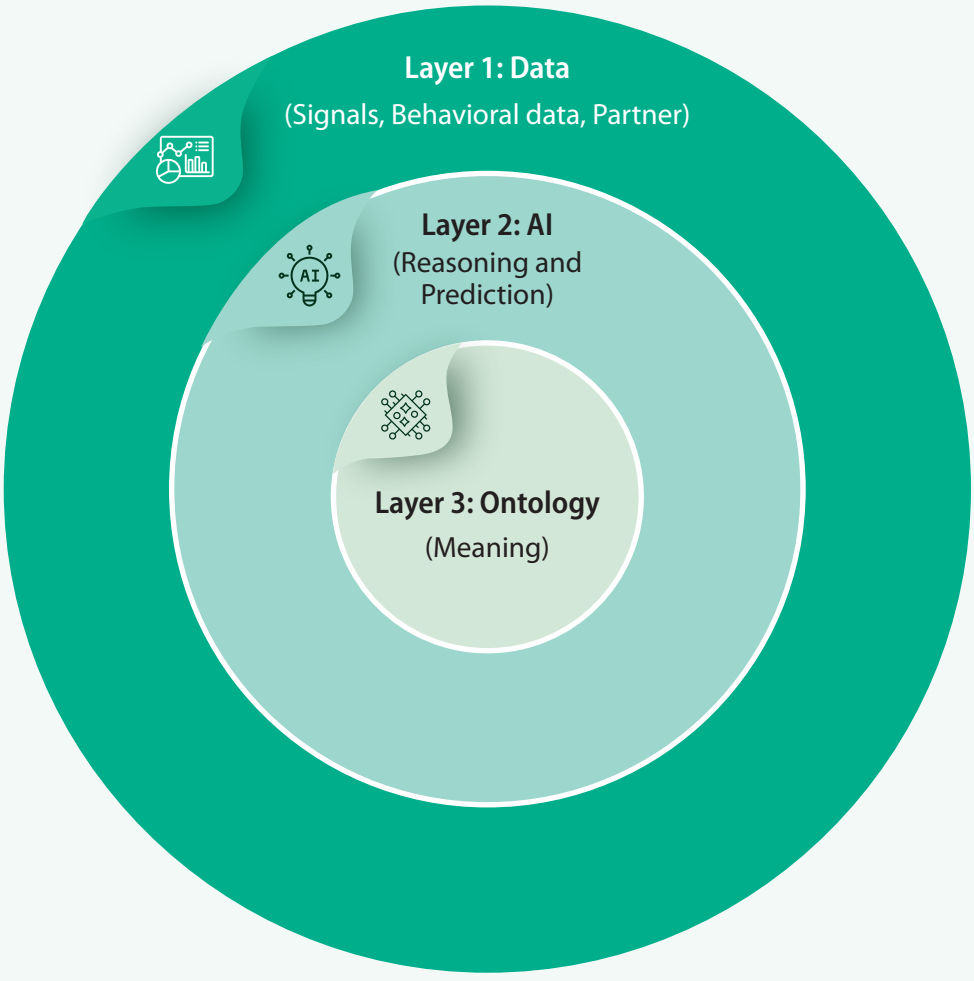

The Five-stage Ontology Journey

For an enterprise marketing organization, the journey from grappling with raw marketing data or fragmented analytics to a clean, operational ontology layer moves through five stages, each building on the last.

Stage 1: Define: The entry point requires marketing and business stakeholders to articulate the essential entities of their operation: What is a persona, precisely, in this organization? What constitutes a journey stage? What does intent mean, and how does it differ from engagement? What buying roles are relevant to the purchase decision?

These are deceptively difficult questions. Organizations that can answer them with precision are ready to operationalize ontology. Those that cannot will find that AI amplifies their ambiguity rather than resolves it. Typical outputs at this stage include entity definitions for Persona, Journey Stage, Intent, and Buying Role.

Stage 2: Identify: With entities defined, the next stage maps the conceptual relationships between them. Customers belong to accounts. Personas influence buying decisions. Engagement signals indicate intent at certain journey stages but not others. Job titles correlate to buying roles. These relationships form the conceptual skeleton of the knowledge graph, which is the logic that allows AI to reason across entities rather than treating them as isolated data points.

Stage 3: Prepare: Feature engineering translates abstract relationships into data-level representations. Page views are mapped as Engagement Events. Job titles are classified as Persona Classes. Content asset types are tagged to journey stages. This stage bridges the gap between the conceptual model and the data reality, ensuring the ontology is grounded in signals the system actually has access to.

Stage 4: Develop: The Formalization Layer creates the shared contract, that is, the defined classes and relationships that consuming applications will use. At this stage, the ontology becomes explicit and machine-readable: customers, campaign, content, and persona are defined as classes; relationships such as *Engages_with, Converts_to, and Influences* are formalized with precise meaning.

This shared contract ensures that every AI tool in the stack is reasoning from the same semantic foundation.

Stage 5: Deploy: The ontology is deployed into a semantic layer — a Knowledge Graph that AI systems, analytics platforms, and CDPs can read from and enrich over time. At this point, the system is live. An AI querying the graph does not just retrieve a data point; it retrieves a data point with its full relational context.

Knowledge Graph Mechanics and Stack Integration

The Knowledge Graph is the operational artefact produced by the five-stage process. It encodes marketing entities such as personas, accounts, journey stages, and content assets as nodes, and the relationships between them as typed, directed edges. These connections capture meaningful business relationships. For example, a persona may engage with a content asset or influence a buying decision. Because each relationship has both meaning and direction, AI can understand not only that two entities are connected, but how they are connected and why that connection matters.

A node might represent a specific Persona class. A persona node might be linked to a journey stage node to show where that persona currently sits in the buying process, and to a content-type node to show what formats have historically moved that persona closer to a decision.

This relational structure is what separates a Knowledge Graph from a data warehouse. A data warehouse answers the question: what happened? A Knowledge Graph answers: what does it mean, and what should happen next?

For AI implementation teams, the graph must be designed for navigation, and moving through relationships, not just retrieving records. Graph query languages such as SPARQL or Cypher are the appropriate tools for this. Exposing only entity data without the relationship context loses most of the graph's value.

Integration occurs at four primary touchpoints in a typical enterprise marketing stack:

- **CDPs** feed behavioral data into the graph and receive enriched audience classifications in return.
- **Lead scoring platforms** use the graph to score prospects based on buying role and journey stage, not just engagement activity.
- **Journey orchestration tools** route prospects based on what the graph knows about them, rather than just static rules.
- **Attribution and analytics platforms** use the graph to link campaign performance to real context: persona, journey stage, and account relationships, enabling more meaningful revenue analysis.

Knowledge Graph Integration Across the Marketing Stack

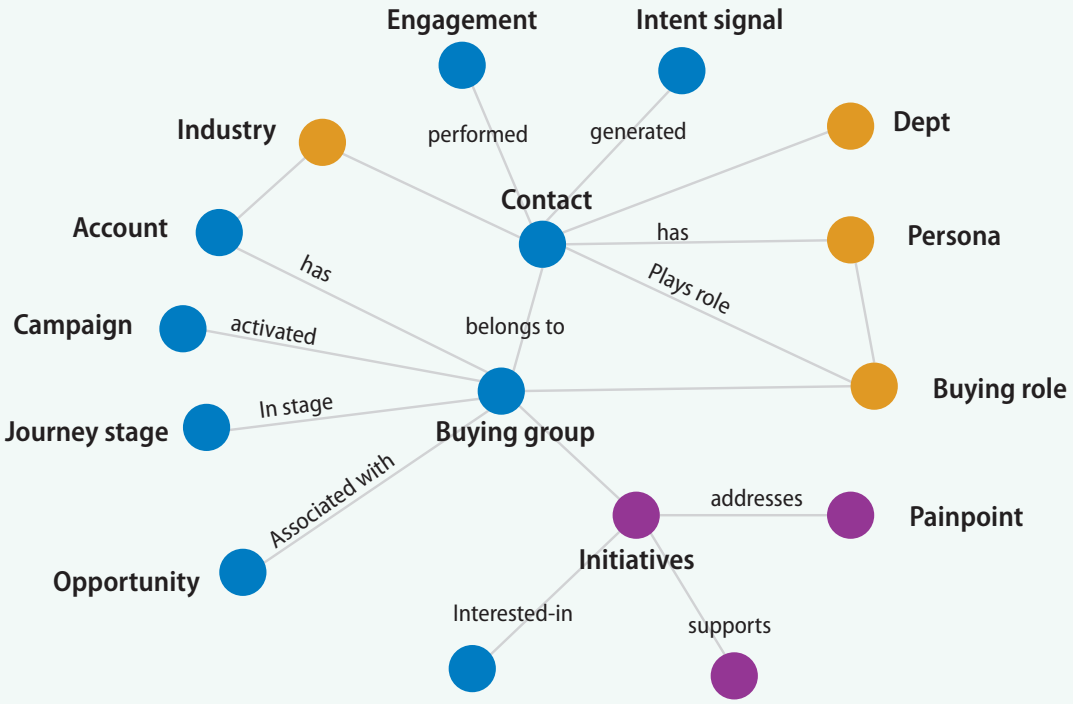

What to Build, Where to Start, and What Ontology Enables

For MarTech leads and AI implementation teams, the hardest and most valuable work is achieving organizational agreement on what core marketing entities actually mean in the specific business context, before any graph schema is written.

The practical entry point follows a recognizable sequence. Begin with one vertical or one product line. Define and agree on the essential entities and formalize the most critical relationships — who your buyers are, what accounts they belong to, and which actions signal where they are in the buying journey. Build a knowledge graph for that scope. Connect your AI tools to it and measure the improvement in decision quality. Once this baseline is proven, expand to new verticals, define new relationship types, and integrate new tools across the stack.

Pre-built, industry-aligned knowledge graphs address the most significant practical barrier to ontology adoption by eliminating the time it takes to build one from scratch. Formalizing the most critical marketing relationships out of the box (Engages_with, Converts_to, Supports, Related_to), aligned to the specific dynamics of verticals such as semiconductors, healthcare, retail, and banking, means implementation teams can begin reasoning contextually without requiring years of graph development or a dedicated data science team to maintain the schema.

Contextual Intelligence is a Competitive Advantage

A Gartner survey of 1,203 data management leaders predicts that through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data.[3]

AI-driven marketing is set to thrive with the context layer added on. Every AI investment a marketing organization makes is only as valuable as the semantic foundation it sits on. Without shared meaning, AI optimizes noise. With it, AI reasons towards superior outcomes.

Ontology has moved beyond being an experimental technology in marketing AI. Early movers who build the foundations now, before contextual intelligence becomes the baseline expectation, will establish an advantage that is difficult to replicate once the market catches up.

Spotlight

Ontology-based Marketing Agent by Infosys Aster

Infosys Aster's marketing agent is built on a robust ontological foundation, bringing together a pre-built knowledge graph, a proven five-stage implementation framework, and deep industry alignment across verticals. For marketing teams ready to move from data abundance to contextual intelligence, the capability is available today.

Infosys Aster ships with pre-established marketing knowledge graphs aligned to verticals, including semiconductors, healthcare, retail, and banking. Teams gain immediate contextual intelligence without requiring deep data science expertise.

What sets Infosys Aster apart

- End-to-end marketing coverage (strategy through measurement)
- Deep AI integration via Infosys Topaz
- MarTech orchestration and data unification across complex enterprise stacks
- Scalability and governance for global B2B and B2C brands.

Infosys Aster Value-driven Ontology Journey Map

Define Ontology Scope

Goal: Define essential and relational vital business processes.

Example: Persona, Journey Stage, Intent, Buying Role.

Feature Engineering

Goal: Define the data, define context, define relationship.

Example: Map page views as Engagement Events. Map job titles as Persona Classes.

Consumption Layer

Goal: Deploy ontology in semantic layer (GraphDB). Data initiated as Knowledge Graph.

AI & Analytic could read from graph. CDP could enrich it.

Define → Identify → Prepare → Develop → Deploy

Identify Key Marketing Concept Map

Goal: Define context of the processes.

Example: Customers belong to accounts, personas influence buying decisions, and engagement signals reveal intent.

Formalization Layer

Goal: Creation of shared contract for consuming applications.

Example: Defined classes – customers, campaign, content and persona. Defined relationships – purchase, engage, etc.

References

1. McKinsey & Company. Gen AI's Broad Reach. 2025.
mckinsey.com/featured-insights/week-in-charts/gen-ais-broad-reach
2. A reintroduction to our Knowledge Graph and knowledge panels.
https://blog.google/products-and-platforms/products/search/about-knowledge-graph-and-knowledge-panels/
3. Lack of AI-Ready Data Puts AI Projects at Risk
https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk

Author

Navin Kumar is an *Industry Principal* at Infosys, specializing in marketing technology, AI-powered marketing transformation, and digital customer experience. With over two decades of experience working with global enterprises, he helps organizations leverage data, AI, and marketing technology to improve customer engagement, marketing effectiveness, and business outcomes.

Navin has led large-scale marketing transformation initiatives across industries, driving innovation in personalization, automation, marketing operations, and AI-enabled decision-making. His work focuses on helping enterprises translate emerging technologies into measurable business value.

For more information, contact askus@infosys.com

Navigate your next