Entity Graph Optimisation

Knowledge Graph Governance for Brand Identity

Santosh Pradhan
MarTech Solutions Architect

SUMMARY

For 20 years, digital marketing optimised one thing: visibility to human eyes on web pages. Every layer built on top — SEO, CMS, analytics, content - existed to serve that primitive. AI has removed human eyes from the discovery loop. The entire document-based stack has no purpose.

GEO and AEO are transitional disciplines that still assume a web of documents underneath.

Entity Graph Optimisation (EGO) is the native architecture for what comes next: a world where brands exist as structured entity definitions in machine-readable knowledge layers — not as pages. EGO is not a content discipline. It is data architecture for brand identity.

0	∞	1
Content Required	Entity Signals	Canonical Definition

TABLE OF CONTENTS

1. THE DISCOVERY COLLAPSE

The web was built for human eyes navigating between documents. Every layer stacked on top — SEO, CMS, CDN, analytics, cookies, martech — exists to serve one primitive: a human reading a page. Remove human eyes from the discovery loop. The entire stack loses its purpose.

✗ DOCUMENT-BASED DISCOVERY	✓ ENTITY-BASED DISCOVERY
User types query. Search engine ranks pages. User clicks. Page loads. Human reads content. Brand is found via keyword density, backlinks, and crawlable HTML. Traffic = sessions on your domain.	User asks AI a question. Model constructs answer from entity graph. Brand is found via entity definition, citation authority, and structured data signals. Traffic = mention frequency in model responses.

61%
Drop in Organic CTR (2025)

2/3
Paid Performance Decline

15k+
Martech Tools Plateau

GEO
Transitional. EGO is Destination

> *" The web will not disappear. It will become infrastructure nobody thinks about — like email protocols. Still running underneath. Invisible to how value moves."*

2. WHAT EGO ACTUALLY IS

Entity Graph Optimisation (EGO) is the discipline of structuring, governing, and distributing an organisation's entity definition across machine-readable knowledge layers, so AI systems can accurately retrieve, reason about, and recommend that entity — without relying on document-based content signals.

How EGO Differs From Existing Disciplines

Discipline	Core Assumption	Primary Output
SEO	Humans discover via search pages	Ranked documents
GEO	AI synthesises from web documents	Cited content
AEO	AI answers direct questions	Structured answers
EGO	AI reasons from entity definitions	Machine-legible identity

The Core Reframe

GEO and AEO are transitional — they still assume a web of documents exists underneath and that optimising those documents is the primary lever.

EGO is post-web. It operates in a world where entities exist in schema stores, not on pages. The question is no longer what should we publish — it is what should we be, inside the model.

When an AI model receives a query, it does not read your blog post. It asks:

→ Is this entity unambiguously defined in sources I trust?

→ Is this entity cited by authoritative sources in my training data?

→ Do structured data signals confirm what this entity claims to be?

→ Is this entity's definition consistent across all surfaces I have seen?

3. THE EGO SIGNAL STACK

GEO practitioners focus on content. EGO practitioners focus on signals. Content is one row in a much larger table. The following signals determine whether an AI system retrieves, trusts, and recommends your entity.

Signal	Nature	Content?	Weight
Entity definition — unique, well-scoped identity	Structural	No	Very High
Structured data — schema.org, JSON-LD, etc.	Structural	No	Very High
Citation network — authority of sources citing	Relational	No	Very High
Co-occurrence — concept association patterns	Semantic	Partially	High
Source authority — presence in model-trusted sources	Reputational	No	High
Factual consistency — same data across all sources	Governance	No	Medium
Entity relationships — adjacency in knowledge	Graph	No	Medium
Content depth — substantive material on domain	Semantic	Yes	Low-mid

> ***"Content is the proof layer — not the signal itself.***
> ***A company claiming expertise with no substantive writing is an unverified entity claim."***

4. THE ENTITY ARCHITECTURE

In the EGO world, a company does not have a website. It has an entity definition — a structured representation of what it is, what it does, and how it relates to adjacent concepts. This definition lives in machine-readable knowledge layers, not in HTML documents.

The Entity Stack

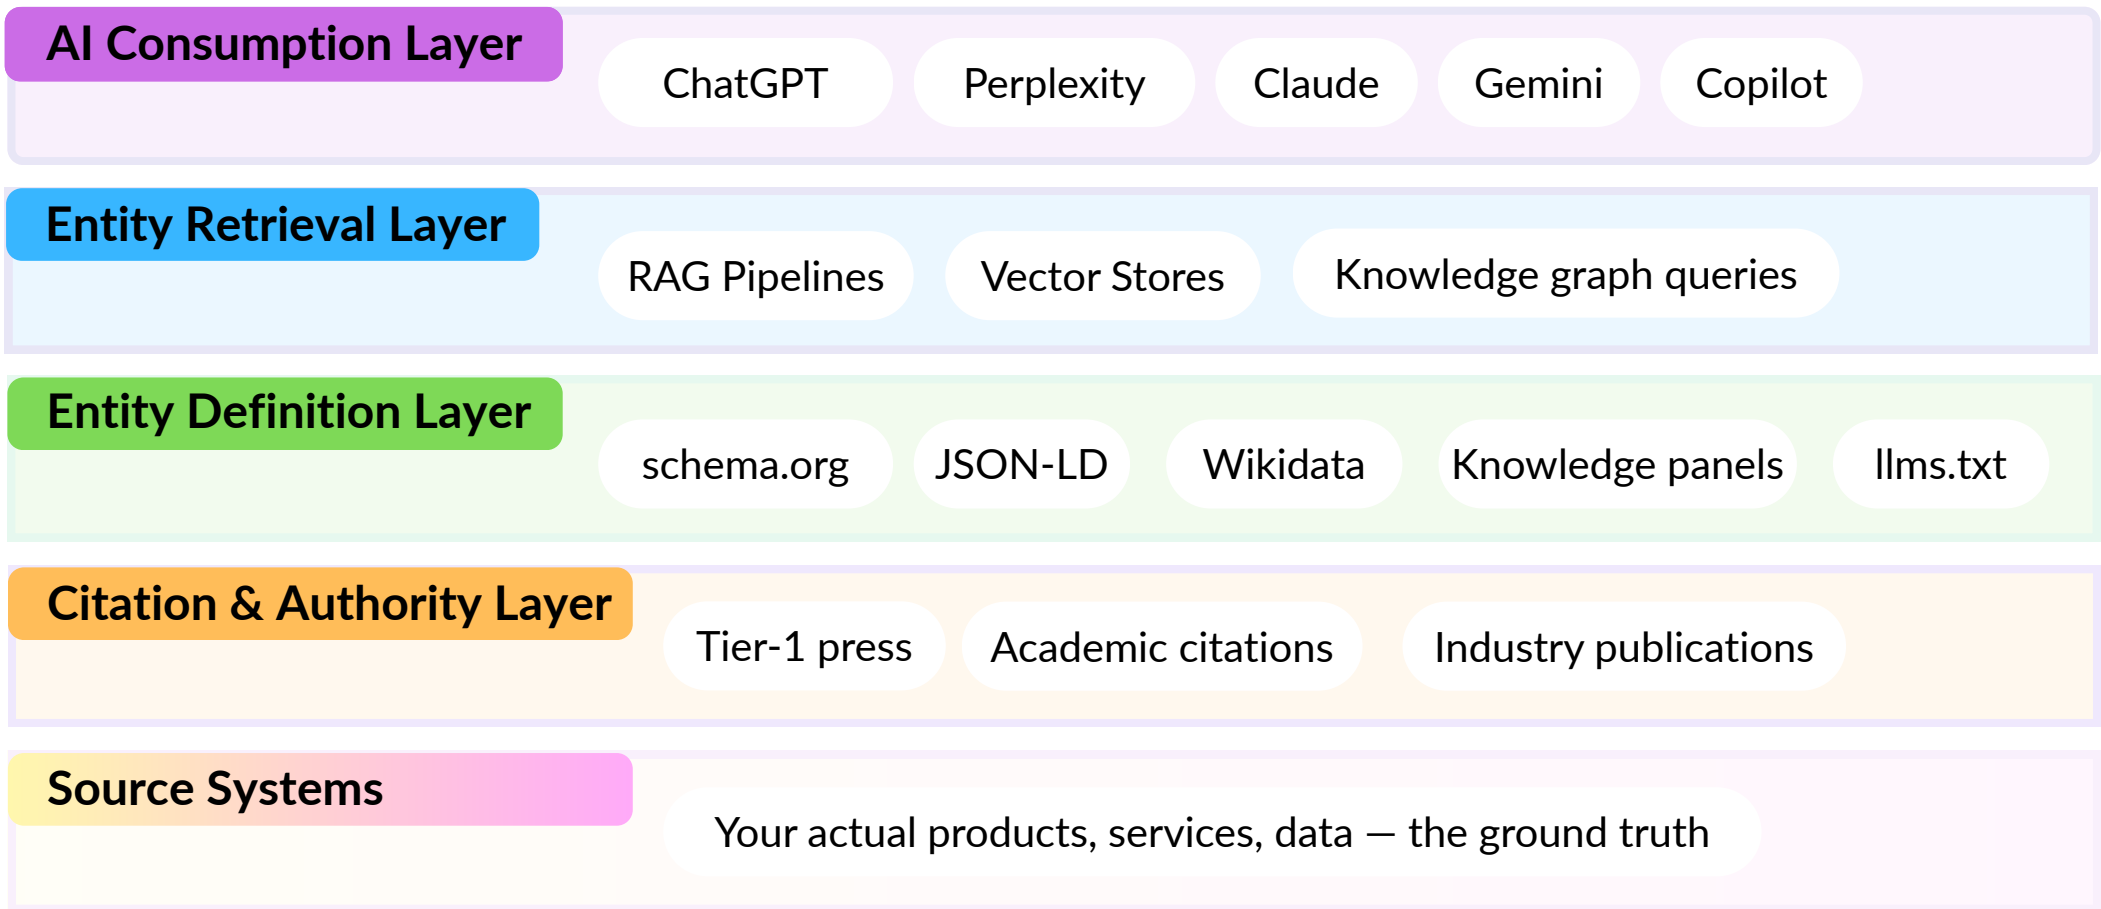

What Replaces the Website

Every company stores itself as a schema in a federated entity store. An AI agent queries that directly. No page load. No bounce rate. No SEO. The entity record defines: identity, products, pricing, trust signals, relationships, and versioned history. Discovery happens inside models — not on pages.

```
Company Entity {
    identity: { name, description, founding_date, HQ }
    products: [ { name, category, use_case, pricing } ]
    trust_signals: { press_citations, certifications, reviews }
    relationships: { competitors, partners, industry_concepts }
    versioned: { history, source_system, confidence_score }
}
```

5. OLD WORLD VS EGO WORLD

EGO does not add a new layer to existing marketing practice. It replaces the foundational assumption. Every function changes.

Old World	EGO World
Content calendar	Entity governance policy
SEO keyword audit	Entity completeness audit
Keyword targeting strategy	Concept ownership map
Backlink building campaigns	Citation engineering programme
Analytics: sessions, CTR, rankings	Metrics: AI mention rate, retrieval frequency
CMO: What is our story?	EGO: How does the model understand us?
Content team at the core	Data architects governing identity
Website as primary asset	Entity definition as primary asset
Publish to be found	Structure to be retrieved
Brand = what you say about yourself	Brand = your definition in a knowledge graph

6. EGO IN PRACTICE

EGO is not a content function. It is a data architecture and identity governance function. The practitioner profile is closer to a solutions architect than a content marketer.

1

Entity Audit

Map your current entity definition across all machine-readable surfaces: Wikipedia, Wikidata, schema.org markup, Google Knowledge Panel, llms.txt, JSON-LD across all pages. Identify gaps, conflicts, and missing signals. This is the EGO equivalent of the technical SEO audit.

2

Citation Engineering

Identify which web sources are cited by AI models for queries in your domain. Engineer presence in those sources through PR, guest contribution, academic reference, and community participation. You cannot be cited if you do not appear in sources the model trusts.

3

Co-occurrence Planting

Associate your entity consistently with owned concepts across trusted surfaces. If you want to own 'field-level timestamp resolution in CDPs' — use that phrase consistently across every authoritative surface until models associate your entity with that concept. This is vocabulary planting, not content marketing.

Governance note:

EGO requires the same factual consistency enforcement as GDPR data accuracy obligations. Conflicting signals across surfaces reduce model confidence and suppress entity retrieval.

7. GETTING STARTED

1

Run an Entity Audit first

Search your brand name in ChatGPT, Perplexity, and Google. What does the model believe you are? Where is the definition incomplete, wrong, or absent? This is your ground zero.

2

Own your Wikidata and Wikipedia entries

These are the highest-authority structured sources models trust. Ensure your entity exists, is accurate, and is consistently described. This single step has outsized impact on model retrieval.

3

Deploy schema.org markup completely

Go beyond basic Org markup. Define products, services, expertise, key people, and relationships to adjacent concepts. LLMs use structured data as a primary source of truth.

4

Engineer citations in trusted sources

Identify which sources appear when AI answers questions in your domain. Get mentioned there — via PR, contributed pieces, research citations, or community participation. Citations are the backlinks of the EGO era.

5

Plant your concepts consistently

Pick the 3-5 concepts you want to own. Use them in the same form across every authoritative surface you touch. Co-occurrence over time is how models associate your entity with a domain. Consistency is the mechanism. Patience is the strategy.

LEARN MORE

https://pradhan.is/ego

GET IN TOUCH

https://www.linkedin.com/in/pradhan-is