

The A.G.E.N.T.I.C. Framework:

Building Brand Authority for AI Search and Agentic Commerce



The A.G.E.N.T.I.C. Framework is a seven-stage method for making a brand visible, credible, and purchasable inside AI-driven discovery, the place where systems such as ChatGPT, Google AI, and Perplexity decide which brands to mention, recommend, and buy from on a shopper's behalf.

01

The framework shifts the objective away from ranking web pages and toward earning semantic authority, the machine-readable trust a model needs before it will name a brand in response to a buyer's question. The seven stages are meant to run in order: Audit, Graph, Equip, Network, Track, Influence, and Convert.

02

Shahzad Safri developed the framework after the traditional SEO playbook lost its edge. A top-of-page ranking means little when the AI does the reading, pulls from a handful of sources, and answers the user without a single click. The brands that come out ahead in this next stage are not those with the largest backlink profiles. They are the ones a machine can parse, verify, and trust enough to recommend and transact with.

The 45-Second Summary


Here is the framework reduced to the parts that carry the most weight:

For more content like and follow me:



@bertblevins

INCGPT.COM



What Does A.G.E.N.T.I.C. Stand For?

A.G.E.N.T.I.C. is an acronym for seven ordered stages, each one solving a specific task an AI platform has to complete before it will recommend you. Skip the foundational tasks and the later ones deliver only a fraction of their potential.

Phase	Letter	What it does
Audit	A	Diagnose where the brand currently stands in AI answers and set a citation baseline.
Graph	G	Build the entity relationships that give the brand a machine-readable identity.
Equip	E	Structure content so AI systems can extract and cite it without friction.
Network	N	Wire inventory and checkout into the transaction layers AI agents use.
Track	T	Measure citations, share of voice, and agent-driven conversions.
Influence	I	Confirm brand claims through trusted external signals.
Convert	C	Capture revenue through agent-to-agent transactions.

Every stage is a separate technical job rather than a marketing slogan. A brand can put one in place on its own and notice a modest gain. The real, compounding advantage comes from following the sequence, where each stage reinforces the next.



Why Traditional SEO Falls Short for AI Search

Traditional SEO is still necessary, but it is no longer sufficient, because AI engines weigh different signals when deciding who to cite than search engines once used to decide who to rank. Classic SEO continues to feed the index, yet it no longer wins the answer.

Ranking tunes a page for a list a human scrolls through. AI tunes for a single synthesized response. The model reads the candidate sources, works out who the brand actually is, checks whether the claims hold up elsewhere, and names a small number of providers at most. If it cannot parse your identity or confirm your claims, you stay invisible, no matter how strong your Google position is.

Three things changed beneath the surface:



How the Seven Stages Work Together

Each stage builds on the one before it, so the value compounds rather than simply adding up. Here is what takes place inside each one.

The Foundation: Audit, Graph, Equip

- 01

Audit.

Put the high-value questions your buyers ask to ChatGPT, Claude, Perplexity, and Google AI Overviews, then record where you are cited, where rivals are named instead, and which content the models pass over. This baseline shapes every later decision.
- 02

Graph.

Construct the entity architecture that tells AI systems who you are. Spelling out explicit facts, such as which person leads which company and how products or services connect, cuts down hallucinations because the model cites your structured data instead of inventing plausible-sounding details.
- 03

Equip.

Reshape content for citation. Open each page with a self-contained answer of roughly 40 to 60 words, then add question-style headers, structured FAQs, bullet points, and clear comparison tables, so the model can lift a clean, quotable response.



The Activation: Network and Track

01

Network.

Connect your catalog, pricing, and checkout to the protocols AI agents rely on to transact. This is the plumbing that lets an agent pull live inventory and finish a purchase.

02

Track.

Measure how often you are cited, how often competitors displace you, how products surface inside AI shopping agents, and the conversion value of agent-referred traffic. Standard web analytics report none of this.

The Expansion: Influence and Convert

01

Influence.

Earn the outside corroboration AI models use to verify your claims. Respected publications, authoritative databases, and review platforms turn your internal story into one the machine can confirm.

02

Convert.

Refine the handoff moment when an agent passes a buyer to your system. Context-aware pricing and clean transaction logic turn a recommendation into revenue.

Although the stages run in order, Track functions as the ongoing feedback loop. A baseline is set during the Audit stage, but Track is established as dedicated infrastructure to measure downstream conversions once the foundation is live.

How A.G.E.N.T.I.C. Prepares a Brand for Agentic Commerce

It wires a brand's catalog, pricing, and checkout into the protocols AI agents use to transact, so an agent can discover a product, recommend it, and complete the purchase without a human ever opening the site. This is the piece most visibility playbooks ignore, and it is the direction the next stage of e-commerce is heading.

01

The guiding rule is to build for the protocol layer rather than any single product sitting on top of it. Products come and go; the standards underneath them endure. The Agentic Commerce Protocol (ACP), maintained by OpenAI and Stripe, is the open standard defining how AI agents and businesses close a purchase, with the brand remaining the merchant of record. It lets ChatGPT take in a merchant's structured catalog, read inventory, and surface products in context.

02

Google built a parallel stack and pushed it into shopping at scale. Its Universal Commerce Protocol (UCP) links merchant catalogs to agent-driven buying, and the Universal Cart, introduced at Google I/O 2026, lets a shopper add items and check out across Search, Gemini, YouTube, and Gmail from one cart.



03

Two more standards round out the stack. Google's Agent Payments Protocol (AP2) handles agent-initiated payments with user-set limits on brand, product, and spend, while its Agent2Agent protocol lets agents talk to one another and coordinate. The Model Context Protocol (MCP), an open standard for connecting AI systems to outside data and tools, then serves as the universal connector between agents and a brand's systems.

04

The framework splits this into two distinct jobs. Network lays the pipes, exposing catalog, pricing, and commerce functions so agents can reach them through these protocols. Convert closes the sale, refining the moment an agent decides whether to recommend and transact. The standards are still maturing, and the consumer products built on them will keep shifting, so the protocol layer is the durable bet. Brands that wire into it now hold a first-mover position as those standards reach full maturity.

Why External Corroboration Decides AI Recommendations

AI models lean on consensus to avoid hallucinating, so a brand with flawless on-site data can still lose the recommendation if nothing outside the brand confirms it. Perfect first-party content is necessary, but on its own it is rarely enough.

01

The logic is simple. The Graph stage tells the AI who you are; the Influence stage tells the AI that you matter. When a model cross-references a claim and finds consistent, positive corroboration across sources it trusts, it recommends the brand with confidence. When it finds silence, it stays neutral, and neutral means unnamed.

02

That corroboration comes from a few places: placement in publications and databases AI retrieval systems treat as high authority, presence on review platforms the models weigh for recommendation confidence, and genuine activity in the communities where AI systems retrieve and learn.

03

An honest limit matters here. You can widen the surface area of accurate, positive mentions. You cannot control what any model ingests, and anyone promising to plant words inside training data is selling a fiction.

04

The retrieval surface is also broader than most brands assume. Yahoo Scout, an AI answer engine launched in January 2026 and powered by Anthropic's Claude model, draws on the open web alongside a knowledge graph of more than a billion entities.

How to Tell If Your Brand Is Invisible to AI

You have an AI visibility problem if an AI engine cannot name you, describe you accurately, or recommend you for the questions your buyers actually ask.

Work through this checklist:

- ☐ When you ask ChatGPT or Perplexity for the best provider in your category, are you named at all?



- ☐ Does the AI describe your products and people accurately, or invent details that are close but wrong?
- ☐ Does a competitor get cited for the exact queries you should own?
- ☐ Can an AI agent pull your live pricing and inventory, or does your catalog stop at a human checkout?
- ☐ Is your authority backed by sources beyond your own website, or does the story begin and end on your domain?

If most of these answers don't work in your favor, you do not have a content problem you can patch with a single blog post. You have a foundation that has not been built yet, and it comes before every tactic.

Which Implementation Path Fits Your Brand?

Most brands enter the framework at one of three points, depending on where they stand today. You do not have to run all seven stages at once.



Frequently Asked Questions

- 01 Is A.G.E.N.T.I.C. just SEO with a new name?**
No. SEO tunes a page to rank in a list a human scrolls. A.G.E.N.T.I.C. tunes a brand to be understood, verified, and named inside an AI-generated answer, and to transact with AI agents directly. The signals, the metrics, and the endgame are all different.
- 02 Is this the same as GEO or AE0?**
Generative engine optimization and answer engine optimization are pieces of it, living mostly inside the Equip and Influence stages. A.G.E.N.T.I.C. is broader. It adds entity architecture, commerce-protocol readiness, measurement, and the agent-to-brand handoff that pure content optimization leaves out.



03

Do I need to run all seven stages?

No. The framework offers different entry paths based on your starting point. Many brands begin with the foundation stages and add commerce or influence work once their source of truth is verified and being cited accurately.



04

How is agentic commerce different from regular e-commerce?

In regular e-commerce, a human browses and checks out. In agentic commerce, an AI agent researches, compares, and can complete the purchase on the user's behalf, often without anyone visiting your site. That requires your catalog and checkout to be machine-accessible, which is the job of the Network and Convert stages.

The Bottom Line

Discovery has moved inside the machine. The next customer may never see a list of ten blue links. That is what A.G.E.N.T.I.C. is built to secure. Audit what the machine sees. Build an identity (Graph) it can read. Equip content it can quote. Network into the commerce layer it transacts on. Track what is real. Influence the consensus forming around the brand. Convert the handoff into revenue.

In AI-mediated discovery, the brands that win are not the ones ranking highest. They are the ones the machine trusts enough to recommend and transact with.