



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

PLAYBOOK

How to Make Sure AI Agents Recommend and Choose Your Business

By Scott Tischler · July 30, 2026 · 9 min read

For twenty years, the game was ranking. You wanted to be the blue link a human clicked. In 2026, a new customer has walked into the room, and it does not click, scroll, or scan a page of ten options. It reads, reasons, shortlists, and increasingly, it buys. That customer is an **AI agent** acting on behalf of a real person, and it is quietly rewriting the rules of how businesses get chosen.

I run Alrecommend.ai, and the single question I get most often from operators right now is some version of: "When someone asks ChatGPT or their agent for the best option, why isn't it me?" This is the playbook I give them. It is not theory. It is what actually moves the needle when the buyer is a machine.

What Does It Mean for an AI Agent to "Choose" Your Business?

An AI agent is software that acts on a person's behalf — it browses, compares, decides, and in a growing number of cases, completes the purchase. Industry estimates in 2026 put roughly **39% of US consumers already using AI for some part of their shopping**, and tokenized agent-payment rails like Mastercard's Agent Pay are launching to let agents transact directly. The agent is no longer a demo. It is a buyer.

Here is the uncomfortable part. When a human searches, they see a menu and exercise judgment. When an agent searches, it often returns **one answer, or a shortlist of three**. There is no page two. There is no scrolling to find you at position seven. You are either in the consideration set the agent assembles, or you are invisible — and invisible to an agent means invisible to the customer behind it.

So the objective shifts. The old goal was traffic. The new goal is **citation and selection**: being the entity the agent retrieves, trusts, and recommends. This is the heart of Answer Engine Optimization, and agentic commerce raises the stakes because the agent doesn't just mention you — it can act.

How Do AI Agents Actually Decide Who to Recommend?

Agents don't "feel" brand affinity. They assemble a decision from signals they can parse. Across the systems I work with, the same factors keep surfacing. An agent recommends the business that is **easy to understand, easy to verify, and easy to trust** — in machine terms.

Break that into what an agent is actually checking:

- **Can I identify this entity unambiguously?** A clear, consistent identity across the web.
- **Can I read the facts I need without guessing?** Structured, machine-readable data — price, availability, service area, hours, specifications.
- **Do independent sources corroborate this?** Reviews, mentions, citations, and third-party validation.
- **Is it present where I retrieve?** The specific sources and indexes the agent pulls from.
- **Can I complete the task?** Clean paths to book, buy, or contact, with no dead ends.

Miss any one of these and you drop out of the shortlist quietly. The agent won't tell you why. It just picks someone else.

Which Signals Make an Agent Pick You? A Priority Table

Not every signal carries equal weight, and operators waste money optimizing the wrong ones. Here is how I rank them for a business starting today.

Signal	What the agent reads	Priority	Typical effort
Structured data	Schema markup: Product, Service, Organization, FAQ, Review, LocalBusiness	Critical	Low–Medium
Entity authority	Consistent identity across the web, knowledge panels, authoritative mentions	Critical	Medium–High
Reviews & reputation	Volume, recency, rating, sentiment, third-party corroboration	High	Ongoing
Machine-readable facts	Clean pricing, availability, specs, service area, hours	High	Low–Medium
Retrieval presence	Being in the sources agents pull from (directories, marketplaces, key indexes)	High	Medium
Transactional readiness	Clear booking/checkout paths, and increasingly agent-payment compatibility	Rising	Medium–High

The pattern to notice: the highest-priority items are also the most controllable. You do not need to beg an algorithm. You need to make your business legible.

How Do You Build Entity Authority an Agent Will Trust?

Entity authority is the degree to which AI systems recognize your business as a distinct, credible, well-corroborated thing in the world. It is the closest thing agentic search has to reputation, and it is the moat most businesses ignore.

Start with consistency. Your name, category, location, and core facts must match everywhere an agent might look — your site, your listings, your profiles, the databases behind knowledge panels. Contradictions create doubt, and a doubtful agent defaults to the safer-looking competitor.

Then build corroboration. Agents weight **independent confirmation** heavily because it is hard to fake. A claim on your own website is an assertion. The same fact confirmed by a review platform, an industry directory, a news mention, and a customer's own words becomes a fact the agent will repeat. This is why earned mentions and genuine reputation work matter more in the agent era, not less. You are not writing for a human skimmer anymore. You are assembling a body of evidence.

One rule I hold clients to: **never manufacture the evidence.** Fake reviews, invented credentials, and inflated claims are the fastest way to get filtered by systems that cross-check sources. The businesses winning in 2026 are the ones whose real strengths are simply made legible.

What Does "Machine-Readable" Really Require?

If entity authority is the reputation layer, **structured data is the plumbing.** An agent that can read your facts in a structured format will trust and cite you over a competitor whose information is trapped in prose, images, or a PDF.

Concretely, this means:

- **Implement schema markup** for the entities that describe your business — Organization, Product or Service, LocalBusiness, Review, FAQ. This is how you hand an agent your facts pre-parsed.
- **Expose the decision-critical fields** an agent needs to compare you: price, availability, service area, specifications, hours, and what you actually do. If a human would need this to choose, an agent needs it structured.
- **Keep it current.** Stale availability or an old price is worse than no data — it makes the agent distrust the whole record.
- **Watch the emerging rails.** The **Model Context Protocol (MCP)**, an open standard connecting agents to tools and data, is seeing rising adoption in 2026. As it matures, the businesses that

expose clean, well-structured interfaces will be the ones agents can actually transact with. Get your data house in order now and you will be ready when the connection layer standardizes.

Think of it this way: every ambiguity you leave in your public information is a small tax the agent has to pay to consider you. Enough tax, and it stops paying.

Where Do Agents Retrieve — and Are You There?

An agent can only recommend from what it can reach. This is the most overlooked part of the playbook. You can have perfect schema and glowing reviews, but if you are absent from the sources a given agent pulls from, none of it counts.

Retrieval presence means showing up in the **directories, marketplaces, review platforms, and authoritative indexes** that feed AI systems in your category. It also means recognizing the **two-channel reality** of 2026: you still optimize for traditional search, because Google AI Overviews now appear in roughly a quarter of searches and pull from the open web — and you separately optimize for the answer engines and agents that retrieve differently. These are not the same channel, and a strategy built for one will underperform on the other.

Vertical matters too. Industry reporting in 2026 consistently shows **domain-specific agents outperforming general models** in their niche. If specialized agents serve your category, being present and legible in that vertical's data sources is worth more than broad visibility everywhere else.

What Should You Do This Quarter?

Here is the sequence I give operators who want to start now, in order of return.

1. **Audit your legibility.** Ask ChatGPT, Perplexity, and Gemini about your category and see if you appear, and what they say. That is your baseline — and often a wake-up call.
2. **Fix identity consistency** across every place your business is described. Kill the contradictions first.
3. **Implement structured data** for your core entities and the fields that drive a buying decision.
4. **Make reputation a system, not a favor.** Steady, real reviews and earned mentions, maintained continuously.
5. **Map and enter your retrieval sources** — the specific directories, platforms, and indexes agents use in your vertical.
6. **Prepare for transaction.** Clean booking and checkout paths today; agent-payment compatibility as the rails firm up.

None of this is exotic. It is disciplined, and discipline is exactly what most of your competitors lack right now. The window where being early is a genuine advantage is open in 2026. It will not stay open. The businesses that make themselves the obvious, verifiable, transactable choice for an agent this year are the ones those agents will keep recommending for years — because trust, once earned by a machine, compounds.

The buyer changed. The work is to be understood by it.

Key takeaways

- ② AI agents return a shortlist of one to three, not a page of ten — you are either in the consideration set or invisible to the customer behind the agent.
- ② Agents choose the business that is easiest to identify, verify, and trust in machine terms: clear identity, structured data, corroborated reputation, and retrieval presence.
- ② Structured data (schema markup) is the plumbing; every ambiguity in your public information is a tax the agent must pay to consider you.
- ② Entity authority — distinct, credible, independently corroborated identity — is the moat, and it can never be faked without getting filtered.
- ② Optimize for the two-channel reality: traditional search (AI Overviews in ~25% of searches) and answer engines/agents retrieve differently and both matter.
- ② Being early in 2026 is a real advantage that will close; trust earned by a machine compounds, so the businesses that become legible now keep getting recommended.

Frequently asked questions

What does it mean for an AI agent to recommend a business?

An AI agent is software that browses, compares, and often buys on a person's behalf. When it recommends a business, it places that business in the short consideration set it returns to the user — frequently just one to three options with no second page. Being recommended by the agent is effectively being chosen by the customer behind it.

How do AI agents decide which business to choose?

Agents assemble a decision from signals they can parse: whether they can identify the business unambiguously, read its facts in machine-readable form, verify it through independent sources, reach it in the places they retrieve

from, and complete the task. A business that is easy to understand, easy to verify, and easy to trust in machine terms wins the shortlist.

What is entity authority and why does it matter for AI recommendations?

Entity authority is the degree to which AI systems recognize your business as a distinct, credible, well-corroborated thing in the world. It functions as reputation in agentic search — the more consistently your identity appears and the more independent sources confirm your facts, the more confidently an agent will cite and recommend you. It is the hardest signal to fake and therefore the most durable advantage.

Does structured data really affect whether AI agents pick my business?

Yes, decisively. Structured data such as schema markup hands agents your facts pre-parsed, so they can read your price, availability, and services without guessing. A business whose information is trapped in prose, images, or PDFs forces the agent to pay a comprehension tax, and agents routinely default to the competitor whose data is cleaner and current.

What is the "two-channel reality" of AI search in 2026?

It means businesses must optimize for both traditional search and AI answer engines because they retrieve information differently. Google AI Overviews appear in roughly a quarter of searches and pull from the open web, while dedicated answer engines and agents draw from their own sources. A strategy built for only one channel will underperform on the other.

Can I fake reviews or credentials to get recommended by AI?

No — it is the fastest way to get filtered out. AI systems cross-check sources and weight independent corroboration precisely because it is hard to fake, so manufactured evidence tends to create contradictions that lower trust. The winning approach in 2026 is making your genuine strengths legible, not inventing strengths you don't have.

What is MCP and why should a business owner care?

The Model Context Protocol (MCP) is an open standard that connects AI agents to tools and data, and its adoption is rising in 2026. As it matures, it becomes a primary way agents read your information and complete transactions. Businesses that expose clean, well-structured data now will be ready to be found and transacted with as the connection layer standardizes.

Where do I start if I want AI agents to recommend my business?

Start by auditing your current legibility — ask ChatGPT, Perplexity, and Gemini about your category and see whether you appear and what they say. Then fix identity consistency across the web, implement structured data for your core facts, build reputation as an ongoing system, and make sure you are present in the directories and indexes agents actually retrieve from. Do these in order of return, beginning with the audit.

About the author. Scott Tischler is the Founder & Chairman of Alrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express,

MetLife, and UBS — and executive study at Wharton, Harvard, Yale, and Oxford, he helps businesses become the ones AI recommends.