

The Agentic Commerce PDP Checklist

Seven steps to make every product detail page (PDP) readable, trustworthy, and recommendable to AI shopping agents.

+138%

y/y AI-driven retail traffic (Adobe Q3 2026)

3x

conversion from answer-engine visitors (Forrester, Q2 2026)

~70%

of AI visits hidden as "Direct" (Dark Agentic Commerce Traffic)

10-20x

content gap on the typical PDP today

☐ 1 • Let the agents in the store

Allowlist agent crawlers at robots.txt, firewall, and CDN. Crawlers don't render JavaScript: reviews, variants, and galleries that only load via JS don't exist. Add llms.txt and a fresh sitemap.

Quick test: spoof GPTBot against a PDP. A 403 means you're invisible.

☐ 2 • Fix identity and canonicalization

One exact title everywhere, model numbers included. GTINs and MPNs populated. One sizing notation. Clean variant matrix. Engines match your offer to a canonical SKU; inconsistency loses the product card.

☐ 3 • Move specs out of prose, into structure

Every spec trapped in a paragraph or baked into an image may never be extracted. schema.org/Product with offers, variants, and images is the minimum; your Universal Commerce Protocol (UCP) and Agentic Commerce Protocol (ACP) feeds should agree with it. Attributes grow 2-10x vs. the search era.

☐ 4 • Write for the questions, not the keywords

Mine service logs, site search, and review Q&A. Ask the engines about your own products, including the negatives. Answer in a Q&A block on the page and in the feed (ChatGPT's q_and_a field is unbounded).

Every unanswered gap is a doorway to a competitor.

☐ 5 • Add the context only you have

Who it's for and who it isn't, occasions, pairings, skill level. Source it from associate training guides, design briefs, and your reviews. Authentic verbosity wins; off-the-shelf AI filler doesn't.

☐ 6 • Keep pricing and availability honest

Price and stock must agree across page, markup, and feed. Promos and strike-throughs go in metadata, not JS-only rendering. A mismatch costs the agent's trust on the next hundred queries.

☐ 7 • Score it, fix the biggest lever, loop

Put a number on every product's agent-readiness (Catalog Score or your own rubric). Rank fixes by score recovery, not gut feel. Re-score, confirm, repeat weekly. Start with your top 20-50 products.

Retailers: start here

Answer engines first (ChatGPT, Gemini, Copilot), then your own site's crawlability and feeds.

Brands: start here

Retail agents already selling you (Rufus, Sparky), then brand.com and the answer engines.