



PLAYBOOK • OCTOBER 2026 EDITION

Shortlisted.

The new e-commerce playbook for when AI does the shopping.

Buyers now ask AI assistants what to buy. This playbook shows how to get your store on AI's shortlist, and how to keep it there.

Best trail running shoe under \$150?

AI

1 · Ridgeline Trail Pro

\$139

2 · Northpeak Grip 4

\$129

3 · Your brand?

Not listed

Seven chapters.

One shortlist to get on.

Read chapters 1 to 3 together as a team. It takes about 20 minutes. Then run the free audit on one product page and work through the 30-day plan.

!	The last 90 days	p. 3
	Agents that buy, checkout on by default, and the first fight over the purchase.	
01	The decision moved	p. 4
	Why the decision happens before the visit, with three simulated answers.	
02	See your store the way an agent does	p. 6
	The agent's view, and six stores that missed the shortlist.	
03	The eight things agents check	p. 8
	223 checks, eight categories, one quick win for each.	
04	Read your score	p. 9
	Tiers, categories, and fixes ranked by impact.	
05	Product facts that win	p. 10
	Five facts every product page must make clear.	
06	The 30-day plan	p. 11
	From first score to a fix that holds, week by week.	
07	Keep it there	p. 12
	Owners, release checks, and when to scan again.	

Commerce leaders

See where the store stands and which fix moves it first.

Technical teams

Trace each finding to the page, template, or data source that owns it.

Agencies

Audit client stores and turn findings into a scoped proposal.

Examples in this playbook are illustrations, not customer results. A worksheet for one product page is on page 13.

The rules changed. In one season.

Four dated events from this summer and autumn. Each one moves the purchase decision further away from your website, and none of them waited for your next planning cycle.

31 Aug 2026

ChatGPT becomes an ad channel

NEW CHANNEL

OpenAI says its ChatGPT ad business reached a \$1B annualized run rate, about 200 days after launch.

8 Sep 2026

Meta launches Muse, an agent that buys

AGENTS BUY

A personal AI agent that can browse, fill in forms, and make purchases for the person, with a check before it pays.

Sep 2026

Shopify switches it on by default

ON BY DEFAULT

For eligible Shopify stores, direct checkout on Meta surfaces, including Muse when available, is activated by default.

20 Sep 2026

Amazon blocks Muse

THE FIGHT

Amazon shuts the agent out of its store. Shopify said yes, Amazon said no. Your store sits in the middle of the fight over who owns the purchase.

+12%

Meta average price per ad, Q2 2026, year over year

\$11.8B

TikTok Shop U.S. GMV, first half 2026, up 103%

~200

Days for ChatGPT ads to reach a \$1B run rate

Paid reach costs more every quarter. A place on the assistant's shortlist costs no media budget, but only stores that agents can read and trust get one.



An agent may already be able to buy from your store.

If you sell on Shopify, check what is switched on today. Nobody on your team had to turn it on.

Shopify admin > Sales channels > Agentic > Meta AI and Muse

The decision moved before the visit.

+393%

AI-referred retail traffic, Q1 2026

+54%

Conversion of AI-referred traffic vs. other traffic

\$1T

U.S. agentic commerce by 2030

Sources: Adobe Analytics, Q1 and May 2026 · McKinsey, 2025.

A buyer asks an assistant for the best option. The assistant answers with one to three names. If your brand is not one of them, the buyer never arrives.

There is no bounce to investigate and no campaign to blame. Analytics records the absence as nothing at all.

An assistant can only recommend what it can reach, read, and trust. An agent can only buy from a store where it finds a way in.

Those are properties of your store. Your team can check them and fix them.

The new path to purchase

1

A buyer asks an assistant

2

The assistant shortlists one to three names

3

An agent checks price, stock, and policies

4

The buyer chooses, or the agent buys

You cannot report on what you cannot see

MEASURE	WEB CHANNEL	AGENT CHANNEL
Impressions	Reported daily	No impression exists. The model answers instead.
Ranking	Tracked per keyword	One shortlist, one to three names, no page two.
Lost sessions	Bounce rate, funnel drop-off	The buyer never arrives, so nothing is logged.
Competitor position	Share of search	Share of answer. Most tools do not measure it.

Not on the list means no click, no session, no sale.

Same question.

Three stores. Three answers.

The assistant only repeats what it can read and check. Here is how one buyer question plays out for three stores that sell almost the same jacket.

Waterproof hiking jacket for women, under €200, delivered by Friday?

STORE A · AGENT READY

Alpine Shell W is a strong match: **€179**, in stock in size M, free delivery by Thursday, and 30-day free returns. 2.5-layer membrane rated 20,000 mm.

RECOMMENDED, WITH FACTS

STORE B · PRICE LOADS LATE

Coastline Rain Jacket may fit, but I could not confirm its current price or stock. Check the store before you order.

MENTIONED, WITH A WARNING

STORE C · CRAWLERS BLOCKED

Not in the answer.
The assistant could not read the store.

NOT LISTED

What made the difference

SIGNAL	STORE A	STORE B	STORE C
AI crawlers allowed	Yes	Yes	Blocked
Price in the server HTML	Yes	No	Unread
Stock per variant	Yes	No	Unread
Product and Offer schema	Yes	Partial	Unread
Delivery and returns as plain facts	Yes	Yes	Unread

Facts get you named. Missing facts get you a warning. A blocked store gets nothing.

Illustrative example. Store and product names are fictional; the assistant answers are simulated.

What a shopper sees is not what an agent gets.

Many agents read the HTML your server sends first and skip JavaScript. A price that loads late is a price they never see. The same goes for stock and for product data that lives only in scripts.

WHAT A SHOPPER SEES



Summit Trail 2
\$145

● In stock

Add to cart

WHAT THE AGENT RECEIVES

```
<!-- /products/summit-trail-2 -->
<h1>Summit Trail 2</h1>
<span class="price"></span> NO PRICE
<div class="stock"></div> NO STOCK
<!-- no Product JSON-LD --> NO SCHEMA
<script src="/app.js"></script>
```

Three ways stores go missing

01

Blocked

robots.txt rules or bot protection turn AI crawlers away before they read a single page.

02

Unreadable

Price, stock, and specs appear only after JavaScript runs. The agent gets empty tags.

03

Unverifiable

No Product schema and vague claims. The model has nothing it can quote as fact.

5'

The five-minute test

Open one of your best-selling product pages. View the page source (Ctrl+U or ⌘⌘U), not the inspector. Search for the price. If it is not in the source, agents that skip JavaScript will not see it either. Do the same for "in stock" and for "@type": "Product".

Illustrative example. Your report shows the evidence from your own pages.

Six stores.

Six ways to miss the shortlist.

None of these needs a replatform. Each one is a small, specific fix that the right owner can ship. Illustrative cases built from common findings, not customer results.

OUTDOOR GEAR

The forgotten robots rule

```
User-agent: GPTBot
Disallow: /
```

A rule from an old site still turns AI crawlers away. Nobody remembers adding it.

Fix: remove it for the crawlers you choose. [ACCESS](#)

FASHION

The price that loads late

```
<span class="price"></span>
<!-- filled by app.js -->
```

Shoppers see €89. Agents that skip JavaScript see an empty tag.

Fix: render price and stock on the server. [CONTENT EXTRACTION](#)

SKINCARE

One price for every size

```
"name": "Night Serum",
"price": "24.00" // 30 mL and 50 mL
```

The schema shows one price for two sizes. The assistant quotes the wrong one.

Fix: one Offer per variant, with its own price and SKU. [STRUCTURED DATA](#)

FURNITURE

In stock online, sold out in the warehouse

```
"availability": "InStock" // last sync: 9 days ago
```

The agent recommends the sofa. The order fails. Next time, it trusts the store less.

Fix: feed availability from the same source as checkout. [AGENTIC COMMERCE](#)

COOKWARE

“Chef-grade” and nothing else

```
<p>Chef-grade. Built to last.</p>
```

No material, weight, or warranty. The assistant has nothing to compare, so it picks a rival.

Fix: state material, weight, oven limit, and warranty; show proof a buyer can check. [ANSWER READINESS](#)

ELECTRONICS

Returns hidden in a pop-up

```
<button onClick="openPolicy()">See policy</button>
```

The return terms exist only inside a script-driven modal. The agent cannot answer “can I return it?”

Fix: link a plain returns page; state window and cost on the product. [ANSWER READINESS](#)

6

small fixes, six owners, and no new platform. The audit's eight categories cover each of these areas, so the report points your team to the right one.

223 checks.

Eight things agents need.

The audit tests each step, from your robots rules to your checkout signals. Every check runs on every scan, free plan included. The method is public, and the scanner behind it, Agent Lighthouse, is open source. Here is what each category asks, and one quick win to start with.

01

Access & Crawl Control

Can AI crawlers reach your pages?

QUICK WIN

Review robots.txt and bot protection for the AI crawlers you want, such as GPTBot, OAI-SearchBot, ClaudeBot, and PerplexityBot.

02

Content Extraction

Can an agent read your product content without running your JavaScript?

QUICK WIN

Send price, stock, and key specs in the first HTML response.

03

Machine Discovery

Sitemaps, feeds, and llms.txt that tell agents what you sell and where.

QUICK WIN

Keep the sitemap current with every live product URL. Add an llms.txt.

04

Structured Data

Product, offer, and organisation schema that models quote as fact.

QUICK WIN

Add Product and Offer JSON-LD with price, currency, and availability.

05

Answer Readiness

Copy an assistant can lift into an answer: specs, prices, policies.

QUICK WIN

Write sizes, materials, delivery, and returns as plain facts, not slogans.

06

Agent Interfaces

MCP, OpenAPI, and well-known endpoints that let agents act, not only read.

QUICK WIN

List what an agent can do on your store, then expose the first action.

07

Agentic Commerce

UCP and ACP signals that let an agent check stock and buy.

QUICK WIN

Check which protocols your platform supports and what it needs to switch them on.

08

Agent Operability & Safety

Errors, speed, and safe defaults when an agent works through your site.

QUICK WIN

Return clear status codes and fast responses on product and cart pages.

Get a score.

Fix what matters first.

Below 50
Not Ready

50-69
Needs Work

70-89
Partially Ready

90+
Agent Ready



EXAMPLE REPORT

Needs Work

223 checks · 8 categories

Access & Crawl Control	88
Content Extraction	41
Machine Discovery	70
Structured Data	55
Answer Readiness	64
Agent Interfaces	20
Agentic Commerce	35
Agent Operability & Safety	81

TOP FIXES, BY IMPACT

HIGH

Product pages ship no Product schema

No JSON-LD Product or Offer found on the scanned page.

HIGH

Prices render only after JavaScript runs

The server HTML has no price; agents that skip JS see none.

MEDIUM

robots.txt blocks GPTBot

Disallow: / under User-agent: GPTBot

How to read it in four steps

01

Check the scope

Note the URL and the time. One page is not the whole store.

02

Start with High

Fixes come ranked by impact. Take the top one first.

03

Confirm the evidence

Each failed check shows what the scan found on your page.

04

Name one owner

Every finding says what to change. Give it a name and a date.

The score measures readiness. It does not count AI recommendations or prove sales. Use separate tests for those. Example values are illustrative.

Agents compare facts.

Give them facts to compare.

An assistant that builds a shortlist needs facts it can check. Make these five clear on every product page, in the visible copy and in the structured data, and make sure they agree.

FACT	WHAT GOOD LOOKS LIKE	COMMON MISS
Title	Brand, product name, and the key attribute in one line.	A campaign name or slogan instead of the product.
Variant	Each size and colour has its own clear name and identifier.	The page shows one price for every variant.
Price	Current price and currency in the server HTML and in the Offer schema.	The price loads after JavaScript, or the schema shows an old price.
Stock	Availability for the selected variant, in text and schema.	"In stock" on the page, nothing in the data.
Policies	Delivery time, cost, and return terms one link away.	Policies hidden in a footer PDF or a pop-up.

BEFORE

~~"Premium quality. Built to last."~~

A claim the model cannot check. It has nothing to quote and nothing to compare.

AFTER

Material Recycled mesh upper

Weight 280 g (US 9)

Warranty 2 years

Returns 30 days, free

The minimum Product schema

```
{
  "@type": "Product",
  "name": "Summit Trail 2",
  "offers": {
    "@type": "Offer",
    "price": "145.00",
    "priceCurrency": "USD",
    "availability": "https://schema.org/InStock"
  }
}
```

Fix disagreements at their source: the system that supplies the fact, not the page that shows it.

Thirty days.

From first score to a fix.

Week

1

Get a baseline

Pick three pages that matter: the home page, one best-selling product, and one category. Run the free audit on each. Save the reports. Do the five-minute test from chapter 2.

Owner: commerce lead **Done when:** three scores and their tiers are recorded

Week

2

Fix the top finding

Take the highest-impact fix. Confirm the evidence on the page. Give it one owner and a test that proves the fix, for example “the price appears in the page source”.

Owner: technical team **Done when:** the fix is live and the test passes

Week

3

Make the product facts clear

Work through the five facts in chapter 5 for your top products. Replace one vague claim with a fact on each page. Check that page, schema, and feed agree.

Owner: merchandising and content **Done when:** the five facts are clear on the top products

Week

4

Scan again and lock it in

Rescan the same pages under the same conditions. Compare the categories, not only the total. Add the audit to the release checklist so the next change does not undo the work.

Owner: commerce lead **Done when:** the new score is recorded and the next fix is planned



Example: Needs Work to Partially Ready

One High fix shipped. Prices now render in the server HTML.
Illustrative values; your results depend on your store.

One audit is a snapshot. Keep watching.

Releases, theme changes, and new bot rules can undo a fix without anyone noticing. Readiness is not a project. It is a property of the store that your team keeps.

Scan again after

- ☐ A theme or template release
- ☐ A new app, plugin, or tag manager change
- ☐ A robots.txt, CDN, or bot protection change
- ☐ A platform upgrade or migration
- ☐ A price, catalogue, or feed change at scale
- ☐ Peak season preparation

Add to the release checklist

- ☐ Price and stock appear in the page source
- ☐ Product schema validates on a sample page
- ☐ robots.txt still allows the crawlers you chose
- ☐ Earlier fixes still hold
- ☐ New findings have an owner
- ☐ The score is recorded with the date

One report. Every owner.

ROLE	OWNS	LOOKS AT
Commerce lead	The score, the priorities, and the plan	Total score and tier
Engineering	Server HTML, schema, robots rules, interfaces	Access, Extraction, Interfaces
Merchandising	Titles, variants, specs, and claims	Answer Readiness, Structured Data
Operations	Prices, stock, delivery, and returns data	Agentic Commerce, Operability

Scheduled audits and score alerts catch the drop before an assistant does.

One product page. One owner per fix.

Page URL	_____	Score	_____
Date and time	_____	Tier	_____

CATEGORY	SCORE	OK	FIX	OWNER
Access & Crawl Control		☐	☐	
Content Extraction		☐	☐	
Machine Discovery		☐	☐	
Structured Data		☐	☐	
Answer Readiness		☐	☐	
Agent Interfaces		☐	☐	
Agentic Commerce		☐	☐	
Agent Operability & Safety		☐	☐	

Top three fixes

#	FINDING	OWNER	DONE WHEN
1			
2			
3			

Five-minute test

- ☐ Price is in the page source
- ☐ Stock is in the page source
- ☐ Product schema is present

Next scan

Date	_____
Score	_____



AGENTIC STOREFRONT BY FORKPOINT

See what AI agents see. Then fix it.

The shortlist forms every time a buyer asks, with or without you. Enter one URL. Get a score, the evidence behind each finding, and the fixes ranked by impact.

Run the Free Audit →

Get Expert Help

agenticstorefront.com/assessment · No account for your score · No card

Sample report

agenticstorefront.com/resources/sample-readiness-report/

Scoring method

agenticstorefront.com/resources/readiness-methodology/

Guides

forkpoint.com/ecommerce-is-changing/

Sources: Adobe Analytics, Q1 and May 2026 · McKinsey, 2025 · OpenAI, Aug 2026 · Meta, Sep 2026 and Q2 2026 results · Shopify Help Center, Oct 2026 · Amazon via Retail Gazette, Sep 2026 · Momentum Works, H1 2026. Examples are illustrations.