

SHOPIFY PLUS · STOREFRONT MCP

The Agent Pre-Flight Checklist

8 things to test before you let an AI agent search your storefront.

An AI agent can already search your Shopify Plus store over MCP. The question is which store it sees: the one you merchandised, or a raw catalog dump. Run these eight checks before you trust an agent in front of a shopper. The last one, the parity check, is the one that catches silent merchandising drift.

8 CHECKS

~30 MINUTES

RE-RUN QUARTERLY

BEFORE YOU START

Before you start

You will need: your storefront MCP endpoint, your storefront access token, and access to your ranking rules and Search Instructions. Set aside about 30 minutes. Bias your test queries toward your highest-volume searches, because a fix on a head query moves real revenue.

CHECKLIST

The 8 checks

Check each box as you confirm it. Checks 1 and 2 are a security gate; checks 3 to 7 are the merchandising gate.

- 1. Endpoint and auth.**
Confirm your storefront MCP endpoint resolves and your storefront access token authenticates as a bearer token. Confirm your management server (write access) is on OAuth and is not exposed to any shopper-facing agent.
Pass / Fail: _____
- 2. Read-only verification.**
Confirm no write or mutate tool is reachable on the storefront server. An agent should be able to read your catalog and never change your store.
Pass / Fail: _____
- 3. Pin test.**
Pin a known product to position 1 for a query. Run that query through the agent's search tool. Confirm the pin holds in the agent's result.
Pass / Fail: _____
- 4. Boost test.**
Promote a brand or attribute. Run the query through the agent path. Confirm the lift shows up in the agent's results.
Pass / Fail: _____
- 5. Promotion test.**
With a scheduled promotion live, confirm the promoted items surface to the agent the same way they do on your storefront.
Pass / Fail: _____
- 6. Out-of-stock test.**
Confirm sold-out items are deprioritized or flagged in agent results, not ranked first.
Pass / Fail: _____

7. Parity test LOAD-BEARING

Run the same five queries through your storefront search box and through the agent's `search-text` tool. Compare the two result sets side by side. Look for merchandising drift: does a pinned product hold the same position, do promoted and demoted items keep their lift, do your Search-Instruction effects (synonym expansion, intent handling, facet ordering) show up the same way. Where the agent set diverges from your storefront, the agent is reading around your ranker for that query.

#	Query	Storefront top results	Agent top results	Drift?
1	-----	-----	-----	<input type="checkbox"/>
2	-----	-----	-----	<input type="checkbox"/>
3	-----	-----	-----	<input type="checkbox"/>
4	-----	-----	-----	<input type="checkbox"/>
5	-----	-----	-----	<input type="checkbox"/>

8. Feedback loop test.

Confirm Agentic Feedback is wired in your `llms.txt` and a test report routes through the Shopify app proxy, so agents can flag a discovery mismatch back to you.

Pass / Fail: -----

Reading your results

GATE

Checks 1 and 2 fail: stop. Fix auth and read-only enforcement before any agent touches the store. This is a security gate, not a merchandising one.

DRIFT

Any of checks 3 to 6 fail, or any drift in check 7: the agent is reading around your ranker for those queries. Your pins, boosts, and promotions are not carrying through to what the agent surfaces.

PASS

Everything passes: the agent is reading through your engine, and the store it shows a shopper is the store you merchandised. Re-run quarterly and after every launch.

The fix for drift is the same in every case. Route agent search and browse through the same ranking engine your storefront uses, so the agent inherits the merchandising you already configured rather than a raw relevance list.

Want us to run this live on your catalog?

We will take your own queries, run them through your storefront and through an agent path on the same engine, and show you exactly where they match and where they drift. 30 minutes, your terms, your products.

[Book a demo →](#)