

AI & Agentic Discovery Readiness

Assistants and agents increasingly summarise, compare and shortlist before a customer ever reaches a site. This framework assesses whether a retailer is legible to them.

What changes

When a model answers rather than lists, a brand competes on whether its proposition, range, pricing and credibility can be extracted and stated accurately by something that never renders the page. Presence is not the same as legibility.

The assessment areas

Entity clarity (is the brand unambiguously identifiable), product and attribute legibility, structured data and feed integrity, third-party corroboration, retrieval-friendly content depth, and the accuracy of what assistants currently say about the brand.

How it is scored

Each area is scored on evidence gathered from repeatable prompts and public signals, with the exact prompt and response recorded. Gaps are ranked by whether they distort the proposition, omit it, or merely under-serve it.

What it is not

Not an AI content programme, not a tooling recommendation, and not a claim about assistant traffic volumes. It is a readiness view that feeds prioritisation.

In summary

- Six assessment areas covering legibility, corroboration and accuracy
- Evidence captured as recorded prompts and responses
- Gaps ranked by distortion, omission or under-service
- Feeds discovery prioritisation, not a standalone AI project

Important

This framework assesses observable representation. It does not estimate assistant-driven revenue or traffic share.

Where this applies

Specialist Growth Programmes — Discovery and AI readiness work is delivered as a focused specialist workstream.

The Growth Diagnostic — Discovery is one of seven lenses; the diagnostic tells you whether it is the constraint worth funding.