Search for Agents: giving AI agents a search endpoint on your website

27 November 2026 · 5 min read · blogs.iictp.org/cloudflare-search-for-agents-website-search-endpoint

Short answer: Search for Agents is a Cloudflare feature, introduced on 25 November 2026, that indexes a website and exposes a search endpoint on the site's own domain for AI agents. A single setting creates the index, starts crawling and advertises the endpoint in page metadata. It matters because it offers site owners a way to hand agents the right passage instead of having them crawl many pages.

TL;DR The short version

- Search for Agents, introduced by Cloudflare on 25 November 2026, indexes a website and exposes a search endpoint for AI agents on the site's own domain.
- One setting creates the index, starts crawling and advertises the endpoint in page metadata.
- It lets you hand agents the right passage instead of having them crawl many pages; test it on a section of your site and check what it returns.

Read online, listen to it, or change the text size: https://blogs.iictp.org/cloudflare-search-for-agents-website-search-endpoint

What is Search for Agents?

It builds on Cloudflare AI Search, a service that indexes content for retrieval, and on Markdown for Agents, which converts HTML pages to Markdown when a client asks for it. Search for Agents puts index creation and endpoint exposure behind a toggle for a zone. Cloudflare says it is available to all customers, including the free plan with limits.

How it works

- **Index.** Cloudflare creates an AI Search instance with default settings and crawls the zone. Reranking and cache options can be changed.
- **Discovery in HTML.** Pages receive meta tags named agent:instructions, agent:endpoint, agent:method and agent:documentation.
- **Discovery in Markdown.** With Markdown for Agents on, the same details appear in front matter.
- **Query.** An agent calls the .well-known/search path on your domain with a query.
- **Response.** The endpoint returns retrieved content in Markdown, not a generated answer.
- **HTTP QUERY.** It also accepts the QUERY method, with a plain text or JSON body.

What is new compared with crawling and sitemaps

A sitemap tells a crawler what exists. This endpoint answers a specific question with matching passages, on request. The agent keeps responsibility for writing the answer, and the site owner keeps the traffic rules: Cloudflare says access follows your bot management, AI Crawl Control and firewall settings.

What the announcement reports

- The post shows example requests against Cloudflare's own developer documentation, with ranked results and scores.
- Natural language questions work depending on configuration.
- QUERY was published as an IETF Proposed Standard, RFC 10008, in June 2026, which Cloudflare gives as a reason it suits search requests that do not change server state.

Limitations and open questions

- The post presents no data showing that agents already look for these tags at scale, or that the endpoint raises visibility or referrals.
- The meta tag names are shown as Cloudflare's convention; do not assume other agents or search engines treat them as a standard.
- Indexed content can go stale between crawls.
- A search endpoint makes retrieval easier for any permitted agent, including ones you would rather not serve.

Sources were read on 27 November 2026. Products, rules and data change, so check the linked originals before you rely on any detail.

How to try it

1. Check that nothing on the zone is public by accident, because indexing will surface it.
2. Decide which agents may access your content and set bot and crawl controls first.
3. Enable the feature from the zone dashboard or API.
4. Run the questions your customers really ask and note wrong or outdated passages.
5. Watch server and analytics logs for calls to the search path before drawing conclusions.
6. Keep your normal search optimisation in place.

What it means for SEO and content teams

Retrieval quality depends on clear headings, short self-contained passages and one current version of each fact. That work helps people too. Measure agent calls separately from visits, and agree with legal and product owners on what agents may retrieve. Related controls are covered in robots.txt and AI training control.

Frequently asked questions

- **What is Search for Agents?** A Cloudflare option that exposes a search endpoint for AI agents on your own domain.

- **Does it generate answers?** No. It returns retrieved content in Markdown.
- **Is it free?** Cloudflare says it is available to all plans, including Free with limits.
- **Does it improve Google rankings?** The announcement makes no such claim.

Related reading

- robots.txt and AI training control
- Generative AI controls in Google Search
- RAG explained
- AI Practitioner certification

> Related credential: AI Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Cloudflare. Introducing Search for Agents. 25 November 2026. blog.cloudflare.com/search-for-agents

Published by IICTP, an independent technical credential body, on 27 November 2026. The current version of this post is at https://blogs.iictp.org/cloudflare-search-for-agents-website-search-endpoint. More guides: blogs.iictp.org.