

MCP SERVER

NO CODE

CLOUD HOSTED

Firecrawl MCP for AI Agents

Get clean Markdown from any website for your research and data extraction.

Firecrawl turns any website into clean Markdown for your AI agents. It handles the messy parts of web scraping like JavaScript rendering and anti-bot protections, so your agent gets structured data instead of a wall of HTML. It's built for high-quality data extraction at scale.

A+ Quality Score 100/100

markdown-conversion

llm-data

web-crawling

data-extraction

structured-data

dynamic-content



The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest
managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.

— Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT`https://edge.vinkius.com/{token}/mcp`

Claude



Cursor



VS Code



Windsurf



Grok



Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1)
Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

Get started at cloud.vinkius.com — connect your AI agent in under 60 seconds.

Firecrawl MCP

4 tools available

Cloud-hosted on Vinkius

When you need your AI to know what's happening on the web, you usually run into a wall of messy HTML, pop-ups, and bot blockers. It's frustrating to build a pipeline that breaks every time a site updates its layout or adds a new cookie banner. This MCP fixes that by handling the heavy lifting of web navigation and content extraction. Instead of your agent struggling to parse a page, it gets a clean Markdown file that it can actually understand and use immediately. You can grab a single article, crawl an entire documentation site to build a knowledge base, or search the web for specific facts and get the results back in a structured format. It's a huge shift from manually scraping data to having your agent just know the web. If you're looking for a reliable way to feed live data into your projects, finding this in the Vinkius catalog makes it easy to get started without worrying about the underlying infrastructure. You just provide an API key and your agent suddenly has a window into the live internet, ready to do the research or data gathering you need.

Core Capabilities

01 — Scrape pages into Markdown

Extract clean content from any URL while handling JavaScript and pop-ups.

02 — Crawl entire websites

Recursively visit every page on a domain to build a full content index.

03 — Map site architecture

Discover all URLs on a website to understand the structure before scraping.

04 — Search and extract web content

Perform a web search and get the content of the top results in one step.

05 — Bypass bot protections

Automatically handle cookies, banners, and anti-bot measures.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/firecrawl — connect your AI agent in three steps.

- 01 Connect your Firecrawl API key to the MCP.
- 02 Tell your agent which URL or search query to target.
- 03 Receive clean Markdown content or a list of discovered URLs.

The bottom line is you get clean, usable web data without the headache of manual scraping or bot blocking.

Built For

This is for the developer who's tired of writing custom scrapers for every new site or the researcher who needs to turn a hundred blog posts into a structured database.

AI Agent Developer

Building autonomous agents that need to browse the web to find real-time information.

Data Engineer

Creating automated pipelines that turn messy web content into structured data for internal use.

Market Researcher

Gathering competitive intelligence from multiple sources without manual copy-pasting.

What Changes When You Connect

- 01 Stop worrying about anti-bot blocks. The system handles proxies and headers automatically so your agent doesn't get blocked.
- 02 Get data in a format your agent actually likes. It converts messy HTML into clean Markdown, saving you from token waste.

-
- 03 Index entire documentation libraries quickly. You can grab every page on a site to build a complete knowledge base.
 - 04 See the full picture of a website first. You can discover all hidden links so you don't miss important data.
 - 05 Combine searching and scraping into one step. Your agent can find and read the best results in a single turn.
-

Real-World Applications

Building a RAG system for company docs

An engineer asks the agent to index the help center, and it recursively crawls the entire site to gather all the content for the vector database.

Automated news summaries

A user wants a daily summary of a niche topic, so the agent searches the web to find and read the latest articles in one go.

Competitive price monitoring

A researcher asks for the latest prices on a competitor's site, and the agent scrapes the page to get the data in a clean format.

Mapping out a new lead list

A sales lead wants to see all the pages on a target company's site, so the agent maps the site to find every available URL.

Patterns to Avoid

Trying to parse raw HTML

✗ AVOID

Asking the agent to read the HTML and hoping it can find the data manually.

✓ INSTEAD

Use `scrape_page` to get clean Markdown instead of raw code.

Manual scraping for every site

✗ AVOID

Writing a custom Python script for every new URL you need to monitor.

✓ INSTEAD

Use `crawl_site` to handle the recursive work automatically.

Getting blocked by bot detectors

X AVOID

The agent keeps hitting 403 errors when trying to access a site.

✓ INSTEAD

Let Firecrawl handle the proxy and header rotation via `scrape_page`.

The Right Fit

Use this MCP if you need your agent to see the live web, especially if the sites have heavy JavaScript or anti-bot measures. It's the best choice for building RAG systems that need fresh data or for agents that perform web research. Don't use this if you only need to read static, simple text files or if you're looking for a way to interact with a specific private API. If you already have a custom scraping infrastructure that works perfectly for your specific needs, you might not need this, but for almost any other web-based data task, this is the standard.

Firecrawl Web Scraping for Automated Data Extraction

Most people trying to get data off the web end up stuck in a loop of manual work. You find a site, copy the text, deal with weird formatting, and then try to paste it into a spreadsheet or a database. If the site has a lot of pages, you're stuck clicking through every single link one by one.

This MCP changes that by letting your agent do the heavy lifting. It goes out, navigates the site, handles the pop-ups, and brings back clean Markdown. You get a structured dataset instead of a pile of messy HTML.

Firecrawl Markdown Conversion for LLM Training and RAG

Building a knowledge base usually requires a lot of manual cleaning. You have to strip out nav bars, footers, and ads just to get to the actual content.

With this MCP, your agent just asks for the content and gets it. It handles the cleaning automatically so you can focus on building the actual application.

It's a tedious process that eats up hours of engineering time.

Firecrawl 4 Tools for Web Data Extraction

Use these tools to scrape, crawl, map, and search the web to get clean Markdown for your agent.

#	TOOL	DESCRIPTION
01	crawl_site	Recursively visit every page on a domain to build a full content index for your agent.
02	map_site	See every URL on a website to understand the structure before you decide what to extract.
03	scrape_page	Get a single URL's content as clean Markdown while handling all the JS rendering and pop-ups for you.
04	search_web	Perform a Google-like search and get the content of the top results back in one step.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

U Find the latest news on renewable energy from the last 24 hours and summarize the top 3.



Latest Renewable Energy News

SOURCE	HEADLINE	SUMMARY
TechDaily	Solar Breakthrough	New perovskite cells hit 30% efficiency
GreenNews	Wind Farm Expansion	500MW project approved in North Sea
EnergyHub	Battery Storage	Solid-state tech reaches milestone

Summary: The focus remains on high-efficiency solar and large-scale wind infrastructure.

U Look at the documentation for the new API and tell me how to authenticate.



Authentication Guide

To authenticate with the new API, you need to follow these steps:

- 1. **Generate a Key:** Go to the developer dashboard to create your API key.
- 2. **Header Setup:** Include the key in your request header as `X-API-Key`.
- 3. **Example Request:**

```
`` bash
curl -H "X-API-Key: YOUR_KEY" https://api.example.com/v1/data
``
```

Note: Keys expire every 90 days.

- U** Map out the entire blog section of this site so I can see how many posts they have.



Site Map: Blog Section

I found 42 URLs in the blog section:

- /blog/getting-started
- /blog/advanced-tips
- /blog/case-study-1
- /blog/case-study-2
- ... and 38 others.

Total Posts Found: 42

Frequently Asked Questions

01 Can Firecrawl help my agent bypass bot protections?

Yes, it's built to handle those issues automatically so your agent doesn't get blocked by common anti-bot measures.

02 Does Firecrawl convert websites to Markdown?

It does, which makes it much easier for your agent to read the content without the noise of HTML tags and scripts.

03 Can I use Firecrawl to crawl an entire website at once?

Yes, you can use it to recursively visit every page on a domain to build a full content index for your agent.

04 How does Firecrawl handle JavaScript-heavy websites?

It renders the JavaScript first, so your agent sees the actual content as it appears to a human user.

05 Is Firecrawl good for building RAG systems?

It's a great choice because it provides clean, structured data that is perfect for feeding into a vector database.

06 Can Firecrawl search the web for me?

It can perform a web search and then extract the content from the top results in one single step.

07 How does Firecrawl pricing work?

Firecrawl uses a credit-based system. You get 500 free lifetime credits to start (no credit card required). Base cost is 1 credit per page scraped. Advanced features like JSON extraction (+4 credits) or enhanced mode (+4 credits) consume additional credits per page. Paid plans start at \$16/month with 3,000 monthly credits.

08 Can Firecrawl handle JavaScript-heavy websites?

Yes! Firecrawl renders pages in a full browser environment before extracting content — this means it handles React, Next.js, Angular, and any other JavaScript framework. It also automatically bypasses common anti-bot protections, removes cookie consent banners, and waits for dynamic content to load before extraction.

09 What formats does Firecrawl return?

Firecrawl can return content in multiple formats: Markdown (default and most popular for LLM consumption), HTML, raw HTML, structured JSON (with LLM-powered extraction), screenshots, links, and page metadata. You can request multiple formats in a single call.

Go Live in 60 Seconds

Get your connection token from cloud.vinkius.com, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

`https://edge.vinkius.com/[TOKEN]/mcp`

CLIENT	WHERE TO CONFIGURE
 Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 Cursor	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"firecrawl": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
 Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

 Ask ChatGPT

 Ask Claude

 Ask Perplexity

 Ask Gemini

 Ask Grok

READY TO CONNECT

Firecrawl is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Firecrawl. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	July 2026
MCP Server	Firecrawl MCP
Server ID	019d7599-9802-73fd-9e98-eba0bf01c694
Platform	Vinkius Cloud for AI Agents
Endpoint	<code>https://edge.vinkius.com/{token}/mcp</code>

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/mcp/firecrawl.