MCP SERVER | NO CODE | CLOUD HOSTED

ZenRows MCP

Scrape any website like it's local.

ZenRows gives your AI client industrial-grade web scraping capabilities. It lets you bypass anti-bot systems, scrape raw HTML from any site, and pull structured data even from complex JavaScript applications. Use it to collect clean datasets across the entire web.

A+ Quality Score 100/100

headless-browser | proxy-rotation | anti-bot-bypass | data-extraction | html-parsing | web-automation

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/zenrows

SCAN TO VIEW

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.

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.

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.

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	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	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.

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.

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.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken 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.

Agent Activity

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

FinOps

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

Tool Health

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

Firewall & DLP

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

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.

ZenRows MCP

10 tools available
Cloud-hosted on Vinkius

Connect ZenRows through Vinkius to give your AI agent access to the live internet. You don't have to manually manage proxies or fight CAPTCHAs anymore. This MCP lets you scrape websites exactly how a human would, handling everything from raw HTML dumps to complex data structures. Need content from a dynamic Single Page Application (SPA)? Use the headless browser functionality to render JavaScript before grabbing the text. Want clean academic datasets? ZenRows can automatically convert messy webpage layouts into clean Markdown, ready for model training. You simply instruct your agent what you need—whether it's localized product pricing across different countries or structured e-commerce data—and the MCP handles the heavy lifting of proxy rotation and anti-bot evasion in the background.

Core Capabilities

01 — Extracting raw website content

Retrieve the complete HTML code from any target page, regardless of how many layers of protection it has.

02 — Handling dynamic JavaScript sites

Use a headless browser to execute site scripts and capture the final rendered view of complex web applications.

03 — Extracting structured data from tables and lists

Automatically parse common website elements, turning messy text into usable JSON records without needing custom selectors.

04 — Bypassing advanced anti-bot security

Execute scrapes against sites protected by major services like Cloudflare or DataDome using specialized bypass technology.

05 — Targeting content by geography

Run scraping jobs using residential proxies assigned to specific countries for accurate, localized data collection.

One Click on Vinkius — From Prompt to Execution

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

01 First, subscribe to ZenRows and enter your API key into the Vinkius catalog.

02 Next, prompt your AI agent with a natural language request—for example, 'Scrape this page using US proxies' or 'Extract all product names from this site'.

03 Finally, the MCP executes the scrape, handles all anti-bot measures internally, and sends back clean, usable data in the format you requested.

The bottom line is: your AI agent becomes a web data collector that doesn't fail when sites try to stop it.

Built For

Data Engineers and Market Analysts who are tired of writing brittle scraping scripts or copy-pasting incomplete screenshots. If your job requires knowing what's on a website right now, this MCP saves hours.

Market Analyst

Monitors global product pricing and availability across dozens of different countries by running geo-targeted scrapes.

Data Engineer

Builds robust data pipelines that automatically bypass bot detection systems, eliminating the need for constant maintenance scripting.

AI Researcher

Quickly gathers large volumes of clean Markdown text from niche websites to train or fine-tune language models.

What Changes When You Connect

01 Skip the boilerplate code. You don't write complex proxy rotation loops; you just ask your agent to run a scrape_geo job for localized data, and it handles the rest.

02 Get clean text datasets instantly. Use scrape_markdown to convert full articles into structured Markdown, perfect for feeding directly into model training without manual cleanup.

03 Bypass tough security walls. Need data from Cloudflare-protected sites? The scrape_antibot tool gives your agent the necessary bypass power when other methods fail.

04 Handle modern websites easily. If a site loads content with JavaScript (like an app), use scrape_js to force rendering and capture that dynamic content, instead of getting blank HTML.

05 Stop guessing what data you'll get. Instead of manually selecting CSS selectors, run scrape_autoparse, and the MCP extracts structured JSON from common layouts automatically.

Real-World Applications

Monitoring competitor pricing changes

A market analyst needs to know if a rival changed its price in Germany. They instruct their agent to run a scrape_geo job, specifying 'de' for the proxy location and asking it to pull product listings, ensuring accurate regional data.

Collecting academic research text

A researcher needs thousands of clean articles on AI ethics. They use scrape_markdown to collect content from a niche blog, automatically stripping away ads and sidebars so only the core article text remains for their dataset.

Debugging site rendering issues

A developer suspects a page isn't loading correctly. They use get_screenshot to capture a live visual preview of the target URL, instantly verifying if the front-end is broken before writing any code.

Building robust data pipelines

A data engineer needs to scrape multiple pages that require specific elements to load. They use scrape_wait to pause the job until a confirmation message appears on the page, guaranteeing the necessary content is available before extraction.

Patterns to Avoid

Assuming simple API access

✗ AVOID

The user tries to run basic HTTP requests and gets empty HTML because the site loads its data via JavaScript after the initial load.

✓ INSTEAD

You must use scrape_js. This tool runs a full headless browser, simulating a real user, which forces the website to execute all necessary scripts before you extract the content.

Ignoring anti-bot protection

✗ AVOID

The scraping job fails instantly with a 'Cloudflare challenge' error page, resulting in zero data and requiring manual intervention.

✓ INSTEAD

Use scrape_antibot. This tool engages the specialized bypass technology needed to get past major security hurdles like Cloudflare or DataDome without failing.

Copying raw HTML blindly

✗ AVOID

The agent pulls raw HTML using scrape_html, which includes navigation menus, ads, and boilerplate code that pollutes the final dataset.

✓ INSTEAD

Run scrape_markdown. This process cleans up the text automatically, delivering only the main body content in a clean Markdown format for immediate use.

The Right Fit

Use ZenRows if your data source is a public website that requires scraping or complex rendering. You need this when simple API keys aren't available because the required information lives deep within a web page, especially one protected by anti-bot measures. Don't use it if you are trying to scrape content from a service that provides a dedicated JSON endpoint; in that case, a basic data retrieval tool is faster. However, if the site uses JavaScript (like modern dashboards or single-page applications), this MCP is necessary because it handles dynamic rendering via tools like scrape_js.

Getting Data from Websites Used To Be an Absolute Nightmare

Today, gathering data means opening a dozen browser tabs. You click here to get raw HTML; then you have to copy-paste that mess into another tool and write regex just to pull out the product name. If the website changes its layout—even slightly—your entire process breaks, forcing you to start over.

With this MCP, you stop managing proxies or writing brittle selectors. You tell your agent what data point matters, whether it's a price or an author bio, and the system executes the complex scraping job automatically, giving you clean data ready for analysis.

ZenRows Gives You Full Control Over Web Data Extraction

The tedious parts that disappear are proxy management, CAPTCHA solving, and the struggle to differentiate useful content from ad filler. You don't worry about whether Cloudflare is up or down; you just get the data.

You can now treat the web like a reliable API endpoint. This MCP makes your AI agent a true web data architect, pulling exactly what you need with minimal effort.

ZenRows: 10 Web Scraping Tools

These tools allow you to perform highly specific, advanced data extraction tasks, ranging from capturing raw HTML to running geo-targeted scrapes.

#	TOOL	DESCRIPTION
01	`scrape_markdown`	Scrapes content and cleans it, converting messy HTML into clean Markdown format by stripping ads and navigation.
02	`scrape_wait`	Pauses the scraping process until a specific element appears on the page, ensuring all required data has loaded.
03	`scrape_premium`	Runs scrapes using high-anonymity residential proxies for better trust and source accuracy.
04	`scrape_html`	Grabs the raw HTML of a page, automatically using anti-bot proxies to prevent blocks.
05	`scrape_autoparse`	Extracts structured data directly from a page, turning visual layouts into usable JSON objects.
06	`get_screenshot`	Generates a URL that immediately returns a screenshot image of the target webpage for verification.
07	`scrape_geo`	Runs a scrape using a proxy limited to a specific country, which is vital for localized content.
08	`scrape_custom`	Executes advanced scraping jobs that require highly specific parameters and selectors.
09	`scrape_js`	Uses a headless browser to run JavaScript and scrape content from modern, dynamic websites.
10	`scrape_antibot`	Performs scraping with maximum anti-bot bypass settings for highly protected corporate or financial sites.

See It in Action

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

U Scrape 'https://example.com' and return the content in Markdown.

I've retrieved the page content and converted it to clean Markdown. Here is the structured text, with all ads and navigation boilerplate removed for your review.

U Bypass Cloudflare and scrape the rendered HTML of 'https://protected-site.com'.

Cloudflare protection successfully bypassed. I've used a headless browser to render the JavaScript and captured the full HTML state for you. What data points would you like me to analyze from this source?

U Get a screenshot of 'https://news-portal.com/breaking-news'.

Screenshot generated successfully. You can view the fully rendered page here: https://zenrows.com/screenshot/.... Would you like me to extract the text content from this page as well?

Frequently Asked Questions

01 **Can ZenRows handle JavaScript-heavy websites?**

Yes, absolutely. Use the scrape_js tool to activate the headless browser. This runs all the site's scripts and captures the full rendered state, so you get content from dynamic Single Page Applications.

02 **How does ZenRows bypass Cloudflare?**

The MCP handles this using specialized anti-bot technology. If a site blocks requests, calling scrape_antibot activates the necessary bypass measures to successfully retrieve the page content.

03 **Is scraping geo-targeted data simple with ZenRows?**

It's straightforward. Just use the scrape_geo tool and specify the country code (like 'gb'). The MCP automatically routes your request through a residential proxy located in that specific region.

04 **What is the difference between scrape_html and scrape_markdown?**

scrape_html gives you the full, raw code dump. If you want clean text suitable for reading or AI ingestion, use scrape_markdown instead; it removes all the ads and navigation clutter.

05 **Does ZenRows handle rate limiting?**

While it doesn't prevent every block, the underlying proxy rotation and anti-bot tools significantly reduce your chances of being flagged. You can also use scrape_wait to slow down requests between pages.

06 **Can ZenRows bypass Cloudflare or other aggressive bot protections?**

Yes. Using the `scrape_antibot` tool, ZenRows activates specialized technology to bypass Cloudflare, DataDome, and other advanced bot detection systems, ensuring you can access the content you need through chat.

07 **How do I get clean Markdown from a website for my AI agent?**

You can use the `scrape_markdown` tool. ZenRows will retrieve the page and automatically convert the content into structured Markdown, removing ads and navigation headers to provide a clean input for LLMs.

08 **Can I see what a dynamic page looks like before extracting data?**

Absolutely. Use the `get_screenshot` tool to retrieve a direct link to an image file representing the fully rendered target page, helping you verify that JavaScript content is correctly displayed.

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 `"zenrows": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → 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

ZenRows 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 ZenRows. 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	June 2026
MCP Server	ZenRows MCP
Server ID	`019d7628-075c-7109-9cc3-cf353ee24999`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

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/zenrows.