

MCP SERVER

NO CODE

CLOUD HOSTED

# Browse AI MCP for AI Agents

## Automating Structured Data Extraction from Websites

Browse AI automates web data extraction by allowing your AI agent to run complex scraping robots and monitor website changes. You can trigger single-page extractions, perform bulk operations across dozens of URLs, or check for updates on competitor websites—all through natural conversation.

**A+** Quality Score 98.33/100

web-extraction

no-code

data-monitoring

automation-robots

structured-data

web-scraping



# 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Browse AI MCP

10 tools available

Cloud-hosted on Vinkius

Need structured data from a webpage but don't want to write Python scripts? This MCP lets your AI client take full control of web scraping tasks. Instead of manually navigating dashboards and copy-pasting metrics, you simply tell your agent what data you need. It executes the necessary extraction robots on target websites and delivers the final result—structured JSON data—right into your conversation context.

This isn't just a simple fetch; it manages everything from triggering multi-URL runs to checking if specific elements on a page have changed since last week. If you're using Vinkius, you can connect this powerful web intelligence source alongside thousands of others in the catalog, making your agent a true data pipeline.

---

## Core Capabilities

### 01 — Execute Web Scraping Workflows

Trigger specific robots to run tasks on target URLs or initiate multi-URL concurrent extractions.

### 03 — Retrieve Extracted Data and Status

Check the progress of a specific scraping task or instantly download the final extracted JSON data from any successful run.

### 05 — Manage API Usage

Retrieve current credit usage against your monthly plan limits for planning bulk operations.

### 02 — Identify and Check Robots

List all your trained extraction workflows, view their configuration details, or check for the status of running jobs.

### 04 — Monitor Web Changes

List all active web monitoring robots and check if target pages have changed since their last scheduled scan.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/browse-ai-1](https://vinkius.com/mcp/browse-ai-1) — connect your AI agent in three steps.

- 01** First, connect this MCP and provide your Browse AI Secret API Key to your agent.
- 02** Next, tell your agent which robot you want to run or what data you need. Your agent will use the necessary tools to initiate the extraction task on the target URL(s).
- 03** Finally, once the job completes, your agent provides the structured JSON payload directly, allowing you to analyze and act on the live web intelligence.

The bottom line is: your AI client handles all the complex steps of scraping, monitoring, and data retrieval so you don't have to interact with a separate dashboard.

---

## Built For

This MCP serves anyone who needs structured information from the open web. If you spend time copying metrics, comparing competitor sites, or manually checking dashboards for changes, this is built for you.

### Data Analyst

You use it to run bulk extractions across multiple datasets, pulling massive amounts of structured data directly into your agent's memory for immediate analysis.

### SEO Specialist

You execute scraping robots on competitor URLs to generate instant market reports or verify content changes without opening a browser tab.

### Market Researcher

You monitor specific target websites for critical metrics, getting alerted only when the data you care about actually changes.

---

## What Changes When You Connect

- 01** Get real-time data results instantly. After running a job using `get_task_data`, you get the structured JSON directly in your chat, skipping manual downloads.

- 
- 02 Stop wasting time on single pages. Use `run_bulk_task` to run an extraction robot across dozens of URLs simultaneously and download all results with `download_bulk_data`.

---

  - 03 Stay ahead of competitors without constant checking. Use the monitoring tools to watch for specific changes, getting notified only when necessary.

---

  - 04 Know your limits upfront. The `list_credits` tool lets you check API usage before launching a massive data pull, saving you from unexpected overages.

---

  - 05 Simplify complex workflows by grouping related actions. You can list all available workflows with `list_robots` and then trigger them without needing to remember every function name.
- 

---

## Real-World Applications

### Monitoring Competitor Pricing Changes

A marketing manager needs to know if a competitor changed their pricing structure. Instead of manually visiting the product page daily, they connect this MCP and use `list_monitors` to set up an automated check. The AI agent then alerts them immediately when prices deviate from the stored baseline.

### Verifying Large Data Sources for Research

A researcher needs confirmation that a large data source is stable. They check the robot configuration using `get_robot` to verify the fields, then use `run_robot` to test it on a new URL, ensuring the output structure matches their expectations.

### Analyzing Market Trends Across Multiple Sites

A data analyst needs a dataset of product specifications from ten different vendor websites. They use `run_bulk_task` to send all ten URLs to the robot simultaneously and then retrieve the unified results using `download_bulk_data`, which saves hours of individual scraping.

---

# Patterns to Avoid

---

## Treating scraping like single-shot actions

### X AVOID

Running a robot once, getting the data, and then having to repeat the entire process manually when the source page changes.

### ✓ INSTEAD

Set up web change monitors using `'list_monitors'`. This lets your agent watch the site automatically and only retrieve new data when a monitored field actually updates.

---

## Forgetting about rate limits

### X AVOID

Launching multiple massive, concurrent scrapings without knowing how many API calls you've already used, leading to job failure.

### ✓ INSTEAD

Always check your account status first. Use `'list_credits'` before initiating any large-scale data collection or bulk operations.

---

## Mixing up task IDs

### X AVOID

Running a scrape and then forgetting the specific Task ID, making it impossible to retrieve the final JSON payload later.

### ✓ INSTEAD

Always record the task status using `'get_task'` immediately after running. You can then use that confirmed success state to reliably call `'get_task_data'`.

---

## The Right Fit

Use this MCP if your primary need is structured, extracted data from web pages. If you are comparing competitor prices or pulling product specifications in bulk, this is the right tool. Don't use it if you just need to read a single article and summarize its text; those tasks work better with standard reading tools. Furthermore, don't try to scrape content that requires logging into a private user dashboard; this MCP only works on publicly accessible pages. If your goal is pure data visualization based on pre-existing spreadsheets, use a dedicated spreadsheet tool instead.

---

---

## Browse AI for Web Scraping: Automating Competitor Data Collection

Right now, tracking competitor activity means opening three different tabs, navigating to each product page, and manually copying the price, key features, and availability status into a spreadsheet. It's time-consuming, error-prone work that slows down your entire market intelligence process.

With this MCP, you tell your agent which robots to run on competitor URLs. The system handles the repetitive extraction, compiles all the data points from every single page, and delivers one clean, structured JSON payload that you can immediately process.

---

## Browse AI for Data Monitoring: Tracking Dynamic Web Changes

The biggest time drain in market research is constantly checking if a website has updated—a new pricing tier, an altered feature list, or a changed policy. Doing this manually means wasting hours on pages that haven't moved.

By setting up monitoring jobs, your agent watches the page for you. You only get notified and retrieve data when something specific changes, giving you a real-time edge without the manual overhead.

---

# Browse AI: 10 Tools for Web Data Extraction and Monitoring

Use these tools to manage robots, run bulk tasks across many URLs, monitor page changes, and download all extracted web data.

#	TOOL	DESCRIPTION
01	<code>run_bulk_task</code>	Runs an extraction robot across multiple URLs concurrently to collect data from many sources at once.
02	<code>download_bulk_data</code>	Gathers and returns all extracted results into a single unified JSON array after a multi-task run completes.
03	<code>get_bulk_task</code>	Checks the execution status of a large, multi-URL bulk scraping operation.
04	<code>get_robot</code>	Retrieves the detailed setup and configuration for a specific web scraping robot.
05	<code>get_task_data</code>	Downloads the final, structured JSON data payload from a successful web scraping task.
06	<code>get_task</code>	Checks the current status of any specific data extraction task run through Browse AI.
07	<code>list_credits</code>	Retrieves your current usage count and quota limits for API credits.
08	<code>list_monitors</code>	Lists all currently active web monitoring robots that check for page changes on a schedule.
09	<code>list_robots</code>	Lists all your active extraction and monitoring robots configured in Browse AI.
10	<code>run_robot</code>	Triggers a configured robot to extract data from a specified target URL, returning a task ID.

---

## See It in Action

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

**U** Can I pull all product specs from three different vendor sites at once?



You got it. I'll run the bulk extraction now for those URLs.

*Status:* Running 3/3 tasks...

*Progress:* 65% complete.

I've initiated the `run_bulk_task`. Give me a minute, and once done, I'll aggregate all the data into one giant JSON array ready for you to download.

**U** Which of my robots are set up to monitor changes?



I checked your monitoring list. You have two active web monitors:

- **Competitor Price Tracker:** Checks the main pricing page every 4 hours.
- **Product Feature Watcher:** Monitors the 'Specs' tab for key feature changes daily.

The Competitor Price Tracker is currently reporting no change since its last run.

**U** Give me all the extracted data from task t-56d3f after it finished.



Task `t-56d3f` completed successfully. Here are the top 3 items found:

PRODUCT	PRICE (USD)	STATUS
Widget Pro	199.00	Available
Basic Tool	49.50	Low Stock
Elite Model	799.00	Sold Out

I have the full, detailed JSON array stored in context if you want me to filter this list by 'Low Stock' items.

---

## Frequently Asked Questions

---

### 01 Can my AI start a data extraction and let me know when it finishes?

Yes! You can ask your agent to `run_robot` on a certain URL. The agent receives the newly created task ID and can immediately invoke `get_task` sequentially to poll for completion. Once successful, it calls `get_task_data` to dump the JSON right into your conversation without breaking workflow.

### 02 How do bulk operations work natively through the agent?

When you have a trained robot, you can feed the `run_bulk_task` command an array of different URLs. Browse AI's backend splits them virtually into independent task instances. Your agent tracks the bulk run via `get_bulk_task` and eventually pulls the mega JSON merge using `download_bulk_data`.

### 03 Is there a risk of my agent suddenly spending all my API credits?

You have complete visibility—the agent can query `list_credits` at any moment to inform you precisely of your available balance and monthly limit. Also, since tools execute your natural directives, your AI agent only triggers extractions you specifically prompt for.







---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"browse-ai-1": { "url": "..."</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# Browse AI 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](https://cloud.vinkius.com) →

[vinkius.com](https://vinkius.com) · [support@vinkius.com](mailto: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 Browse AI. 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	Browse AI MCP
Server ID	019d7564-459b-71de-ac71-9ac4aabb33f2
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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/browse-ai-1](https://vinkius.com/mcp/browse-ai-1).