MCP SERVER | NO CODE | CLOUD HOSTED

Nimbleway MCP

Web Scraping & Structured Web Search for Your Agent

Nimbleway lets your AI client scrape and search the live web right from your chat window or IDE. Extract content—whether you need raw HTML or clean Markdown—and perform structured searches against major search engines. It's a direct connection to high-performance data collection, letting your agent fetch information without needing external tools or browser extensions.

A+ Quality Score 100/100

data-extraction | proxy-services | web-crawling | stealth-browsing | structured-data | automated-search

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

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.

Nimbleway MCP

10 tools available
Cloud-hosted on Vinkius

Nimbleway connects your AI client directly to web scraping and advanced search capabilities. You can ask your agent to fetch any webpage, rendering it as either raw HTML or clean Markdown format while bypassing typical anti-bot protections. Need market intelligence? Use the MCP to perform structured searches across major search engines and get organized data back immediately. Beyond simple extraction, you manage complex operations: check account usage via `get_account_usage`, monitor active jobs using `list_jobs`, and even inspect proxy configuration details with `list_proxies`. If your AI client is compatible with Vinkius, connecting this MCP gives you access to all these web data tools through a natural conversation. This means complex tasks—like checking pipeline status via `get_pipeline` or viewing account info with `get_me`—become simple commands for your agent.

Core Capabilities

01 — Scrape Web Content

Request the content of any live webpage, receiving it either as raw HTML or clean Markdown text.

02 — Perform Structured Searches

Execute real-time web searches and receive structured data points from major search engines.

03 — Monitor Data Jobs

List, check the status of, or retrieve metadata for active scraping jobs.

04 — Manage Scraping Pipelines

View and inspect details about your configured data extraction workflows.

05 — Check Account Status

Retrieve current bandwidth usage, remaining credits, and overall account metrics.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your specific Nimbleway API Key in your client settings.

02 Your AI agent recognizes the available web data tools (like `extract_markdown` or `search_web`) within Vinkius.

03 You issue a natural language command, and your agent invokes the tool, executing the scrape or search and returning the structured data to you.

The bottom line is that instead of writing code to connect to scraping APIs, you just tell your AI client what web data you need.

Built For

This MCP solves the problem for researchers and developers who hate context switching. It's for anyone whose job requires gathering external information quickly —from competitive analysis to building proof-of-concept agents. You don't need a full development stack; you just need access to live web data via your agent.

Market Researcher

Needs to automate the retrieval of structured search results across multiple industries for comparative analysis, without manually visiting dozens of websites.

Data Engineer

Wants to quickly test extraction rules or check pipeline statuses directly through their chat interface instead of opening a dedicated dashboard.

Software Developer

Requires high-fidelity web scraping capabilities for proof-of-concept agents, integrating live data gathering into the natural flow of writing code and prompts.

What Changes When You Connect

01 Go beyond simple data copying. Use `search_web` to get organized, structured results from live web searches, which is critical for market research.

02 Avoid messy code blocks. Instead of writing complex scraping scripts, simply ask your agent to use `extract_markdown` and receive clean, readable text directly in the chat.

03 Know your limits before you start. Use `get_account_usage` to check bandwidth and remaining credits instantly so you don't hit a paywall mid-project.

04 Manage large operations without leaving your IDE. You can use `list_pipelines` or `list_jobs` to monitor complex data streams right alongside your code development.

05 Bypass common anti-bot measures. The underlying technology handles the tricky rendering, allowing you to focus purely on the data you need from a given URL.

Real-World Applications

Competitive Analysis

A market researcher needs competitor pricing and product features across five different sites. Instead of opening five browser tabs, they ask their agent to use `extract_markdown` on all five pages and then compile the data into a single comparison table.

Building Agent Proofs

A developer is building an automated content generator. They first use `search_web` for current industry trends, then use `extract_html` on the top result to gather raw source material, all within a single session.

Auditing Web Crawl Jobs

A data engineer suspects a pipeline is failing. They immediately ask the agent to run `list_jobs`, find the ID, and then use `get_job` to diagnose exactly why the last extraction failed.

Real-Time Information Gathering

A student needs quick facts about a niche topic. Instead of doing multiple Google searches, they ask their agent to run `search_web`, which returns structured results that are immediately usable in an essay draft.

Patterns to Avoid

Assuming Browser Functionality

✗ AVOID

Thinking you can just copy and paste a URL into your agent and expect clean data without knowing the underlying limitations.

✓ INSTEAD

For reliable, high-fidelity content extraction, use `extract_markdown` or `extract_html`. These tools guarantee that the content is properly scraped regardless of the site's complexity.

Ignoring Usage Limits

✗ AVOID

Running several large scraping jobs back-to-back and only realizing you ran out of bandwidth when your agent fails with a cryptic error.

✓ INSTEAD

Always check your resource availability first by calling `get_account_usage`. This prevents job failure and saves time.

Treating Web Search as General Chat

✗ AVOID

Asking the agent to 'just find me some data' without specifying a format or structure, resulting in unstructured text chunks.

✓ INSTEAD

Use `search_web` specifically. This tool forces structured output and gives you usable snippets, titles, and URLs that are ready for immediate analysis.

The Right Fit

You should use this MCP if your primary need is gathering live, external data—whether it's a webpage's content or organized search results. The strength here is the combination of extraction power (using `extract_markdown` and `extract_html`) and structured searching (`search_web`), all accessible via conversation. Don't use this if your goal is internal processing; for example, if you just need to summarize data already uploaded into a document, stick to standard text generation tools. Also, don't rely on it for persistent scheduling without proper setup—you must check `list_pipelines` first. If you only need to monitor the status of existing jobs, using `get_job` or `list_jobs` is faster than trying to re-crawl everything.

Collecting Web Data Used To Feel Like a Full-Time Job

Today, gathering data means opening multiple tabs, dealing with CAPTCHAs, and constantly switching between the browser, your spreadsheet, and your notes. You copy text from one source, paste it into another, then write down which URL gave you that piece of information. It's slow, tedious, and easily breaks when a website changes its layout.

With this MCP, you tell your agent exactly what you need—a specific web page or a set of search results. The agent handles the complexity of scraping and structuring the data so you just get clean Markdown or structured JSON back in your chat window. You spend zero time on manual extraction.

Web Data Extraction with Nimbleway MCP

The days of manually checking proxies, logging into a separate dashboard just to see if your job finished, or figuring out what bandwidth you have left are over. You don't need to manage multiple platforms.

Now, you treat web data collection like any other command: ask for it, and get the result. It’s instant, conversational, and keeps all your context right where you work.

Nimbleway: 10 Tools for Web Data Access

These tools give your AI client direct access to scraping, searching, and account management features necessary for any data-intensive workflow.

#	TOOL	DESCRIPTION
01	`extract_html`	Pulls the full raw HTML content from a specified webpage URL.
02	`extract_markdown`	Extracts the clean, readable text and formatting (like headers) from a webpage as Markdown.
03	`get_job`	Retrieves specific status details for an individual scraping job run.
04	`get_me`	Pulls general current account information and metadata about your Nimbleway profile.
05	`get_pipeline`	Provides detailed status and configuration information for a specific data scraping pipeline.
06	`get_account_usage`	Checks your account's current bandwidth usage, remaining credits, and overall billing status.
07	`list_jobs`	Generates a list of all completed and active scraping jobs you have run.
08	`list_pipelines`	Lists every available data collection pipeline you've set up in your account.
09	`list_proxies`	Accesses the configuration details for all residential and data center proxies associated with your account.
10	`search_web`	Executes a structured search query against major web engines and returns organized results.

See It in Action

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

U Extract the content of 'https://example.com' as Markdown.

Initializing extraction... I've successfully fetched the page and converted it to Markdown. The content includes headers, lists, and links from the original site. Shall I summarize it for you?

U Search the web for 'latest AI developments 2024' and give me structured results.

Executing web search... I've found several top results regarding generative AI, LLM improvements, and new hardware releases. I've retrieved the snippets, titles, and URLs for each. Would you like to extract the full content of the top result?

U Check my account usage and remaining credits.

Retrieving usage data... You have used 15GB of bandwidth this month out of your 50GB limit. You also have 4,500 successful extraction credits remaining. Your account status is 'ACTIVE'.

Frequently Asked Questions

01 **How does Nimbleway MCP handle website anti-bot measures?**

Nimbleway uses advanced stealth technology to bypass common bot detection systems. This lets your agent scrape protected sites and retrieve content you otherwise couldn't access.

02 **Can I use Nimbleway MCP for general web searching?**

Yes, the `search_web` tool performs structured searches across major search engines, giving you organized snippets, titles, and URLs immediately.

03 **What is the difference between extract_html and extract_markdown?**

Use `extract_markdown` when you want clean, human-readable text (like an article). Use `extract_html` if you need the raw source code structure for development purposes.

04 **How do I check my usage with Nimbleway MCP?**

You call the `get_account_usage` tool. This instantly reports your current bandwidth used, remaining credits, and overall account status so you know exactly how much data you can pull.

05 **Is Nimbleway MCP just for scraping one site?**

No, it's designed for variety. You can use `search_web` for broad research, and then use tools like `extract_markdown` on the specific sites that interest you.

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 `"nimbleway": { "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

Nimbleway 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 Nimbleway. 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	Nimbleway MCP
Server ID	`019d75dd-c98e-71b2-8414-d5033b7ed1ef`
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/nimbleway.