MCP SERVER | NO CODE | CLOUD HOSTED

PhantomBuster MCP

Control web scraping and data extraction via conversation.

PhantomBuster connects web data extraction and automation workflows directly to your AI agent. List all available scraping agents (Phantoms), launch new campaigns with specific arguments, check current usage limits, and retrieve the final extracted leads or structured data—all via natural conversation.

A+ Quality Score 98.33/100

automation | data-extraction | lead-gen | workflow-automation | browser-automation

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

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.

PhantomBuster MCP

10 tools available
Cloud-hosted on Vinkius

Web scraping used to be a clunky process involving dedicated dashboards and manual API calls. Now you can manage your entire lead generation pipeline using only chat commands. Instead of clicking through multiple tabs to see if a workflow is running, simply ask your agent about it. You can check the status of all your defined automation agents or launch an entirely new campaign on demand, supplying custom parameters in conversation. When scraping finishes, you don't have to hunt for the results; you just ask for the latest output and get the clean data ready for immediate use. Because Vinkius hosts this MCP, you access full control over all your web automation tools through one connection point. This means whether you need to audit a configuration or retrieve 10,000 leads from a past run, it's all available inside your agent.

Core Capabilities

01 — View and manage scraping agents

List every defined web scraping tool (Phantom) and every automated workflow you have set up.

02 — Start or stop automation runs

Launch a specific Phantom agent with custom arguments, or immediately halt a runaway process that's consuming credits.

03 — Retrieve extracted data results

Fetch the final output and raw data collected from any completed scraping job for analysis.

04 — Audit automation configurations

Review the precise setup arguments of a Phantom or workflow without running it, ensuring your logic is correct before launch.

05 — Check account usage and limits

Get real-time information on your remaining execution time and overall account settings.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your PhantomBuster API key credentials.

02 Tell your agent what you need, for example: 'List all Phantoms that scrape LinkedIn profiles.'

03 Your agent sends the request using the available tools, returning a list of agents or the requested data directly.

The bottom line is you stop working in web dashboards and start directing your automation through conversation.

Built For

Growth Marketers who get frustrated by manual lead scraping, Sales Operations teams needing rapid CRM enrichment, and Data Engineers who need reliable auditing of complex data pipelines.

Growth Marketing Manager

They launch a new Phantom agent to scrape leads from a specific niche website, then ask the agent to get_phantom_output once the run finishes so they can feed the list into an email campaign.

Sales Operations Analyst

They check the account usage via get_phantombuster_account and audit existing workflows using get_phantom_setup before integrating them into a new sales reporting dashboard.

Data Engineer

When debugging, they list all containers to ensure the correct environment is running, or use abort_phantom if an automated process starts looping endlessly.

What Changes When You Connect

01 You gain immediate visibility into all automation assets. Instead of logging into a separate dashboard, simply ask the agent to list_phantoms or list_workflows to see every tool you own.

02 Stop manual intervention when things go wrong. If a scraping job gets stuck or runs too long, use abort_phantom to halt the process instantly and save credits.

03 Never lose data again. After a successful run, ask for the output using get_phantom_output; the agent pulls the clean results right into your conversation for you to review.

04 Verify complex logic before launching expensive campaigns. Use get_phantom_setup to audit exactly what parameters a Phantom needs without wasting execution time.

05 Track costs and capacity with confidence. Check your remaining credits and overall status using get_phantombuster_account, so you always know if you can afford the next scrape.

Real-World Applications

The Sales team needs a quick list of competitors.

Instead of building a new workflow and hoping it works, the agent first uses get_phantom to check the details of an existing competitor profile Phantom. Once verified, they use launch_phantom to run it immediately with specific industry keywords.

A developer needs to debug data extraction.

The engineer runs a process that fails. They ask the agent to get_workflow details and check the corresponding Phantom configuration using get_phantom_setup, narrowing down whether the error is in the code or the input arguments.

A marketing manager needs historical data.

The manager asks for the output from a run last week. The agent fetches this specific result using get_phantom_output and provides 200 rows of structured lead data instantly, eliminating manual CSV downloads.

A team needs to scale up scraping efforts.

Before launching a massive campaign, the user first asks list_phantoms to ensure all necessary agents are present. They then check get_phantombuster_account to confirm they have enough capacity for the whole batch.

Patterns to Avoid

Trying to run a phantom with incorrect arguments

✗ AVOID

The user launches_phantom and forgets to specify the target URL or search query, resulting in an immediate error and wasted time.

✓ INSTEAD

Before launching, always use get_phantom_setup for that specific agent. This shows you exactly which configuration arguments are required—like 'target_site' or 'search_term'—so your chat prompt is perfect.

Ignoring running processes

✗ AVOID

A scraping job starts but seems stuck, chewing through credits while the user waits for a response that never comes.

✓ INSTEAD

Don't wait. Immediately use list_containers to verify it's still running, and if necessary, use abort_phantom to kill the process cleanly.

Assuming data is available

✗ AVOID

A user assumes that a Phantom run from last week still has the raw data accessible in their current session.

✓ INSTEAD

Always explicitly request historical data using get_phantom_output, specifying which agent and when the data was generated. This guarantees you retrieve the correct dataset.

The Right Fit

Use this MCP if your workflow involves complex web scraping or lead generation where the process requires multiple distinct steps: checking status, configuring parameters, launching, and then retrieving structured output. It is built for automation control.

Don't use it if you simply need to read static information (use a basic text extraction tool) or if your data source isn't web-based (use a database connector). If you just want to list files in a folder, this MCP won't help. You need the dynamic control over running and stopping sophisticated scraping agents; that is what get_phantom, launch_phantom, and abort_phantom provide.

The pain of managing web data extraction today

Right now, if you want to scrape a new dataset, you jump into the PhantomBuster dashboard. You check which agents are active on one screen, then switch tabs to see logs and run status. If something fails or needs adjustment, you have to go back through complex configuration settings just to verify what went wrong. It's a lot of clicking through different dashboards just to get visibility.

With this MCP connected via Vinkius, all that dashboard hopping disappears. You tell your agent exactly what you need—like 'Show me the status of my top 3 scraping agents.' The agent handles the complexity behind the scenes and gives you one clean answer in chat. You get immediate control without ever leaving your conversation window.

Gain total control using PhantomBuster

You no longer have to manually navigate between listing agents, checking account usage (get_phantombuster_account), and finally pulling the results. The agent orchestrates these steps for you based on your natural language command.

This changes everything. You move from being a dashboard operator who clicks buttons to an AI director who simply gives instructions. It's faster, safer, and dramatically more scalable.

PhantomBuster MCP: 10 Tools

These tools let you manage every aspect of web data extraction, from listing available agents to retrieving finalized campaign results.

#	TOOL	DESCRIPTION
01	`abort_phantom`	Stops a web scraping agent that is currently running.
02	`get_phantombuster_account`	Retrieves your account's usage statistics and settings.
03	`get_phantom_setup`	Shows the specific configuration arguments used by a defined Phantom agent.
04	`get_phantom_output`	Fetches the final data and results from a completed scraping run.
05	`get_phantom`	Retrieves detailed information for one specific web scraping agent.
06	`get_workflow`	Gets details about an entire automated workflow setup.
07	`launch_phantom`	Starts a new Phantom execution with specified parameters.
08	`list_containers`	Lists all running and defined web scraping containers.
09	`list_phantoms`	Shows a list of every available Phantom agent you own.
10	`list_workflows`	Lists all the automated workflows configured in your account.

See It in Action

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

U List all Phantoms in my account and show their current status.

I've retrieved 5 Phantoms: 'LinkedIn Search Export' (Success), 'Sales Navigator Search' (Error), 'Web Scraper' (Ready), and 2 others. Which one would you like to manage?

U Launch the 'LinkedIn Search Export' Phantom with ID '12345'.

Phantom 'LinkedIn Search Export' (ID: 12345) has been launched successfully. I'll notify you once the output is ready.

U Show me the extracted data from the last run of Phantom 12345.

I've fetched the results for Phantom 12345. It extracted 150 leads, including names, company titles, and LinkedIn profile URLs. Would you like a detailed breakdown?

Frequently Asked Questions

01 **How do I list all my scraping agents using PhantomBuster MCP?**

You ask your agent to list_phantoms. This tool shows you every defined web scraping agent (Phantom) in your account, letting you see what's available for use.

02 **What if a scrape job is running too long? How do I stop it?**

Use abort_phantom. You just tell the agent to abort_phantom and provide the Phantom name or ID, and it will immediately halt the running process.

03 **Can I see what arguments a specific agent needs before launching it?**

Yes, you use get_phantom_setup. This tool shows you exactly which configuration arguments are mandatory for that Phantom, saving you from errors and wasted credits.

04 **Where do I find the data after running a PhantomBuster automation?**

You retrieve it using get_phantom_output. Just tell your agent to fetch the output for a specific run ID or date, and the clean, extracted data appears right away.

05 **Does this MCP help me track my web scraping costs?**

Absolutely. Use get_phantombuster_account. This tool provides real-time account settings and usage information, so you always know your remaining execution time before launching a job.

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 `"phantombuster": { "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

PhantomBuster 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 PhantomBuster. 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	PhantomBuster MCP
Server ID	`019d75f2-4f03-70ce-b634-5c1d87534276`
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/phantombuster.