MCP SERVER | NO CODE | CLOUD HOSTED

Have I Been Pwned MCP

Audit Your Digital Footprint for Leaks and Breaches

Have I Been Pwned MCP checks if your email or passwords were exposed in known data breaches. It connects your AI agent directly to the trusted HIBP database, allowing you to audit accounts and verify password safety against thousands of historical leaks. Check account involvement or discover details on a specific hack using this MCP.

A+ Quality Score 100/100

data-breach | cybersecurity | identity-protection | password-security | threat-intelligence | account-safety

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/have-i-been-pwned

SCAN TO VIEW

The infrastructure that powers AI agents in the real world.

Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

Have I Been Pwned MCP

5 tools available
Cloud-hosted on Vinkius

This MCP lets your agent act as an instant digital security auditor. You stop guessing if your data is safe and start checking the record. It pulls real-time breach intelligence, verifying whether specific accounts were compromised or if passwords have appeared in public leaks.

Need to check a personal email? Use this MCP to run an account search against major breaches. Worried about old passwords? The system checks for password safety using k-anonymity, meaning your actual password never leaves your client and is always protected.

Beyond checking accounts, you can also use the tool to discover if information has been posted on public paste sites, or explore a full history of major data compromises. This capability puts deep threat intelligence right into your chat window, making complex security auditing simple. By connecting this MCP via Vinkius, you're giving your agent access to one of the internet's most trusted resources for protecting sensitive information.

Core Capabilities

01 — Audit Account Breaches

Checks if a specific email or username appears in any recorded data breach.

02 — Find Public Paste Exposures

Scans public paste sites to see if an account name or email has been leaked there.

03 — Validate Password Safety

Confirms whether a password was ever compromised in a breach without transmitting the full password.

04 — List All Breach Events

Retrieves a comprehensive list of all major data breaches currently tracked by the service.

05 — Get Specific Breach Details

Fetches detailed information about one specific, named data breach event.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/have-i-been-pwned** — connect your AI agent in three steps.

01 First, subscribe to this MCP on Vinkius and obtain your HIBP API Key.

02 Second, input the provided key into your AI client's configuration panel. This authorizes the connection for breach checking.

03 Third, simply ask your agent to 'check if X email was compromised,' or 'is Y password safe?' The MCP runs the query and returns the findings.

The bottom line is you get instant, verifiable data on digital risk without having to visit a separate website or manage API calls manually.

Built For

Anyone dealing with sensitive PII (Personally Identifiable Information) needs this. Security analysts and IT professionals use it daily to check for corporate domain compromises, while privacy advocates rely on its breach history data to advise clients.

Security Analyst

Runs the list_all_breaches tool to track emerging threat vectors or uses search_account_breaches to vet a company's domain integrity after an incident.

Privacy Consultant

Guides clients through checking password safety and running account searches for personal leaks, explaining the risk level of compromised data.

IT Professional

Verifies if corporate user accounts or internal system credentials have been exposed in public breaches using search_account_breaches.

What Changes When You Connect

01 Immediate Risk Assessment: Quickly run account searches using search_account_breaches to see every breach an email has been part of. Stop guessing about your security status.

02 Secure Password Testing: Use check_password_safety to validate if a password was leaked without sending the password itself over the wire. Your data stays protected.

03 Comprehensive Tracking: Access the full history via list_all_breaches and get deep context on any specific event using get_breach_details, keeping you ahead of threat actors.

04 Public Leak Detection: The search_account_pastes tool goes beyond breach databases by checking public paste sites for your leaked credentials or identity details.

05 Single Source of Truth: Instead of hopping between multiple security websites, this MCP consolidates all necessary checks—breaches, pastes, and passwords—in one conversational flow.

Real-World Applications

Vetting a New Client's Security

A consultant needs to advise a client about their overall digital risk. They ask the agent to run search_account_breaches on the client's main corporate email, then use check_password_safety to test several key employee passwords. The MCP returns a clear report of all identified risks.

Investigating an Old Hack

An IT professional remembers a breach from 2016 and wants to know what exactly was compromised. They use get_breach_details, specifying the name of the hack, immediately getting details on data types stolen (passwords, phone numbers, etc.).

Monitoring for Leaked Credentials

A researcher suspects an account might be floating around public forums. They use search_account_pastes to check if the user's email or name has appeared in any publicly accessible paste sites, providing a layer of defense beyond formal breaches.

Building a Risk Report

A security analyst needs to document all potential risks for a client. They start by calling list_all_breaches to get the scope of known threats, then use search_account_breaches on the target account to narrow down relevant exposures.

Patterns to Avoid

Checking only emails

✗ AVOID

A user asks the agent, 'Is my email safe?' and stops there. This only checks account breaches but ignores potential password leaks or public posts.

✓ INSTEAD

To audit completely, use search_account_breaches for the email, then immediately run check_password_safety on multiple strong passwords to cover both accounts and credentials.

Assuming current safety

✗ AVOID

A user thinks because their password hasn't been found in a major breach yet that it is safe forever.

✓ INSTEAD

Run check_password_safety regularly. Even if not listed today, this MCP allows you to verify against the massive growing database of known compromises.

Ignoring public leaks

✗ AVOID

A user only checks formal breach databases and misses data posted on niche forums or paste sites.

✓ INSTEAD

Always use search_account_pastes to catch information that might be leaked outside of major, tracked corporate breaches.

The Right Fit

Use this MCP if your primary need is verifiable threat intelligence regarding compromised credentials and identity data. You must check *what* was breached (search_account_breaches) and *if* a password has been compromised (check_password_safety). Don't use it if you just want to know general industry trends; for that, the list_all_breaches tool is sufficient. If your goal is to manage tickets or update customer records, this MCP is useless—you need a dedicated CRM integration instead. This is purely an intelligence and auditing layer.

The Constant Fear of Digital Compromise

Right now, checking your digital safety feels like detective work. You have to copy-paste emails into one tool, run a password through another service that uses different rules, and then manually cross-reference those results with public paste sites. It's exhausting, time-consuming, and you always feel like you’re missing some crucial piece of data.

With this MCP, the process is conversational. You describe your security concern—like checking if an old password was exposed or if a specific email was involved in a major hack—and your agent executes all necessary checks behind the scenes. It delivers a single, consolidated report telling you exactly what's safe and what needs immediate attention.

Access Breach Data with Have I Been Pwned MCP

The specific manual steps that disappear are the repeated copy-pasting of emails, the need to manually visit different breach databases, and the uncertainty about whether an old password is truly safe. Your agent handles all those calls for you.

Now, auditing your digital life is a direct question you ask your AI client. It's not a multi-step process; it’s just checking account safety using search_account_breaches or verifying credentials with check_password_safety.

Have I Been Pwned: 5 Security Tools

These tools let you run deep security audits, checking for compromised credentials, listing all known hacks, and searching for leaked information on public paste sites.

#	TOOL	DESCRIPTION
01	`check_password_safety`	Checks if a password was ever found in a data breach using k-anonymity, keeping the actual password private.
02	`get_breach_details`	Pulls specific information about one particular recorded data breach by its official name.
03	`list_all_breaches`	Returns a complete catalog of every major data breach event currently documented in the system.
04	`search_account_breaches`	Searches for all known breaches associated with a provided email address or account username.
05	`search_account_pastes`	Checks public paste sites to see if an email or account name has been leaked there.

See It in Action

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

U Has my email 'test@example.com' been involved in any breaches?

Checking breaches for 'test@example.com'... I found 3 breaches: Adobe (2013), LinkedIn (2016), and Canva (2019). Leaked data includes: Email addresses, Passwords, Usernames.

U Is the password 'P@ssword123' safe to use?

⚠️ **DANGER!** This password has appeared in **45,210** data breaches. You should **NEVER** use this password as it is easily guessable and publicly known.

Frequently Asked Questions

01 **How does Have I Been Pwned MCP work with my password?**

It uses k-anonymity when you run check_password_safety. This means the system checks if a password was found in a leak without ever sending your actual, full password to the server.

02 **Can I find out all data breaches with Have I Been Pwned MCP?**

Yes, you use the list_all_breaches tool. This gives you access to a comprehensive catalog of every major breach event recorded by the service.

03 **What is search_account_pastes useful for?**

It searches public paste sites specifically. This finds instances where your email or account may have been posted somewhere outside of formal, tracked data breaches.

04 **Do I need an API Key to use Have I Been Pwned MCP?**

Yes, you must provide a valid HIBP API Key during setup. This key authorizes your AI client to run the security checks against the live database.

05 **Which tool should I use if my email was compromised?**

Start with search_account_breaches. This is the most direct way to see all known breaches linked to that specific account or username.

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 `"have-i-been-pwned": { "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

Have I Been Pwned 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 Have I Been Pwned. 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	Have I Been Pwned MCP
Server ID	`019d8445-c874-716e-8a3f-39896e5f1e63`
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/have-i-been-pwned.