MCP SERVER | NO CODE | CLOUD HOSTED

AbuseIPDB MCP for AI Agents

Auditing IP Risk and Reputation Scores for Network Security

AbuseIPDB MCP instantly audits IP addresses against global, crowdsourced databases. It checks an IP's abuse score, reviews detailed report histories, and maintains a current blacklist of high-risk IPs right from your AI chat client. Stop manually checking security dashboards—get real-time network intelligence effortlessly.

A+ Quality Score 100/100

ip-reputation | cybersecurity | threat-intelligence | network-security | abuse-reporting | data-lookup

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

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.

AbuseIPDB MCP

4 tools available
Cloud-hosted on Vinkius

Need to audit network traffic or vet suspicious IPs? This MCP lets your AI agent manage complex IP reputation checks without you opening a single security dashboard. Instead of digging through multiple reports, you simply ask your agent for the status of an address, and it instantly pulls high-resolution metadata. It's like having a real-time security consultant available in conversation form.

When using this MCP via Vinkius, your AI client takes over the tedious process of cross-referencing data. Your agent can check if an IPv4 or IPv6 address is associated with malicious activity, audit the confidence score to gauge risk likelihood, and even pull detailed reports on past spam or DDoS attempts. It turns massive security data into simple answers, letting you keep your network intelligence verified and precise.

Core Capabilities

01 — Check IP Address Status

Checks an IP address against the AbuseIPDB database to get its reputation score.

02 — Verify Service Operational Status

Confirms if the AbuseIPDB service is currently running and available for queries.

03 — Retrieve Global Blacklist Data

Gets the current list of IP addresses that have been most frequently reported across the globe.

04 — Get Detailed Abuse Reports

Collects a full record and history of reports associated with a specific IP address over time.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and input your AbuseIPDB API Key.

02 Connect your preferred AI client (Claude, Cursor, Windsurf, etc.) through the Vinkius catalog.

03 Ask your agent a natural language question—like 'What is the risk score for 1.2.3.4?'—and get instant results.

The bottom line is you talk to your AI client, and it uses this MCP to pull live security data directly into your conversation window.

Built For

Security Analysts who need rapid threat intelligence. DevOps Engineers struggling with log verification at 2 a.m. Network Administrators facing suspicious IPs that require immediate, deep auditing.

Security Analyst

Checks IP reputations and retrieves official metadata to build incident reports straight from their workflow.

DevOps Engineer

Audits incoming traffic patterns and verifies server logs without having to jump between multiple monitoring dashboards.

Network Administrator

Performs rapid audits on suspicious IPs using natural language prompts, identifying relevant security markers instantly.

What Changes When You Connect

01 Instantly audit any IPv4 or IPv6 address using the `check_ip_address` tool. You get high-resolution reputation metadata right in your chat, eliminating manual dashboard searches.

02 Understand potential threats faster by auditing abuse confidence scores. This feature lets you gauge the likelihood of malicious intent instantly, without guesswork.

03 Build a clear activity timeline for an IP address using `get_ip_abuse_reports`. You can identify patterns of spam or hacking simply by requesting the full report history.

04 Maintain strict network control by querying the global list via `get_abuse_blacklist`. This keeps your system informed about the most currently reported bad actors.

05 Keep your security research flowing smoothly. Use `check_api_status` to ensure the MCP is operational before running any critical, time-sensitive audits.

Real-World Applications

Investigating Suspicious Server Traffic

A DevOps Engineer finds a sudden spike in traffic from an unknown IP. They ask their agent to check the address using `check_ip_address` and immediately see if it has been flagged for past malicious activity, confirming if they need to block it.

Forensic Analysis of Incident Logs

A Security Analyst is reviewing logs from a suspected attack vector. They use `get_ip_abuse_reports` on the source IP to build a timeline, finding evidence of prior spamming or DDoS activity that wasn't obvious in the primary log data.

Implementing New Access Controls

A Network Administrator needs to update firewall rules. They use `get_abuse_blacklist` and cross-reference the top offenders, ensuring all high-risk addresses are blocked organization-wide immediately.

Vetting Partner Connections

An Operations Lead is onboarding a new partner network. Before granting access, they query the IP range using `check_ip_address` to confirm that the entire block has a clean reputation score, minimizing potential security exposure.

Patterns to Avoid

Treating IPs as static data points

✗ AVOID

Simply looking up an IP address once and assuming its risk level remains the same. This fails when an attacker quickly switches source IPs, leaving you with outdated security context.

✓ INSTEAD

Don't just run a single check. Use `get_ip_abuse_reports` to see historical trends and cross-reference that data with the current status from `check_ip_address`. This gives you the full lifecycle picture.

Ignoring systemic failures

✗ AVOID

Running a massive, multi-hour audit without first checking if the underlying service is stable. The whole process halts unexpectedly when the API fails mid-run, wasting time and resources.

✓ INSTEAD

Always start by running `check_api_status`. This simple check confirms your access point is live before you initiate any resource-intensive lookups like querying the blacklist.

Overlooking high-risk IPs

✗ AVOID

Only checking specific IPs mentioned in an incident, and missing broader patterns of attack. The attacker might be using a known botnet IP that you never thought to query.

✓ INSTEAD

Proactively use `get_abuse_blacklist` to get the latest list of most reported addresses. This ensures your team is always aware of current global threat vectors.

The Right Fit

Use this MCP if your core need is real-time, deep IP reputation analysis and historical auditing. If you're trying to validate a single asset for known malice, `check_ip_address` is the perfect starting point. However, don't use it if you only need basic geolocation data—you'll need a different tool for that. Similarly, this MCP doesn't help you write firewall rules; it gives you the intelligence needed to build them. If your goal is merely log parsing without context enrichment, stick to a simple text processor instead.

AbuseIPDB MCP: Auditing IP Reputation Scores for Threat Intelligence

Today, investigating a suspicious connection means juggling multiple dashboards. You copy an IP into one tool for reputation, then switch to another for report history, and maybe a third just to check if the service is even online. It's tedious manual clicking that slows down incident response time.

With this MCP, you simply ask your agent to 'Audit the risk of 203.0.113.4.' Your agent handles all those steps—the status check, the confidence score lookup, and pulling recent reports—and presents a consolidated answer instantly. You get actionable security context without leaving your chat window.

AbuseIPDB MCP: Network Security for Blacklist Management

Manually maintaining network access lists is a nightmare. When a new, massive botnet IP appears in the wild, you have to monitor dozens of feeds and manually update your block list across different systems.

Now, by using this MCP's tools, you can query the current global blacklist with one prompt. This means your team stays ahead of threat actors; they tell you what's currently bad so you can keep your network airtight.

AbuseIPDB: 4 Tools for IP Reputation Auditing

Use these four tools to check an IP's status, retrieve historical reports, and monitor the global blacklist of reported addresses.

#	TOOL	DESCRIPTION
01	`check_ip_address`	Checks a given IP address against the AbuseIPDB database for its current reputation score and metadata.
02	`check_api_status`	Confirms if the AbuseIPDB service is online and ready to accept queries, ensuring your audits run smoothly.
03	`get_abuse_blacklist`	Retrieves the current list of IP addresses that have been reported most frequently globally.
04	`get_ip_abuse_reports`	Gathers a detailed history and list of reports for a specific IP address to trace past activity.

See It in Action

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

U What is the overall risk score for 192.0.2.5? Also, show me its report history.

IP Analysis: 192.0.2.5

- **Abuse Confidence Score:** High (82%)
- **Last Seen Activity:** Spamming bulk emails.
- **Report Count (90 Days):** 67 reports.

Detailed History:

DATE RANGE	ACTIVITY TYPE	SEVERITY
Last Week	Bulk Emailing	Medium
3 Weeks Ago	Port Scanning	Low

Verdict: This IP is actively used for spamming and should be flagged immediately.

U Run a quick check on this range of IPs. Is anything on the global blacklist?

Global Blacklist Check Complete.

We found 3 addresses in your query that are currently listed on the active abuse blacklist. These include:

- `10.255.255.1` : Associated with massive brute-force attacks.
- `203.0.113.88` : Flagged for continuous DDoS activity.

Please review these immediately; they require network isolation.

U Is the AbuseIPDB service currently running? I need to run a big audit.

System Status Check: OK.

The AbuseIPDB API is fully operational. We can proceed with your security research workflow. The connection status shows green across all metrics, so you don't need to worry about downtime when auditing the IP reputation.

Frequently Asked Questions

01 **How does AbuseIPDB MCP help me audit my network traffic?**

It lets your AI agent check any specific IP address against a massive, global database of reported IPs. This gives you instant visibility into whether that IP is associated with known malicious activity or spam.

02 **Can I use AbuseIPDB MCP to find out why an IP was reported?**

Yes. You can retrieve the detailed reporting history for a given IP address. This shows you patterns, like if it's being flagged repeatedly for spamming or hacking over time.

03 **What if I need to know if my whole system is safe right now?**

You can check the current global blacklist using AbuseIPDB MCP. This shows you a list of IPs currently reported by the community, helping you proactively block known bad actors.

04 **Is this better than checking multiple websites manually?**

Absolutely. Instead of clicking through different security sites, your AI agent pulls all the necessary reputation data and historical context into one place for quick review.

05 **How do I verify if AbuseIPDB MCP is working before a big audit?**

The MCP includes a status check tool. Running this first confirms that the service is fully operational, so you can trust the data you get back during critical security investigations.

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 `"abuseipdb": { "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

AbuseIPDB 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 AbuseIPDB. 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	AbuseIPDB MCP
Server ID	`019d8411-8b47-732c-bb99-d9dc914375ff`
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/abuseipdb.