

MCP SERVER

NO CODE

CLOUD HOSTED

IPQualityScore (IPQS) MCP

Audit user inputs instantly: IPs, URLs, Emails.

IPQualityScore (IPQS) delivers real-time security checks for IPs, emails, URLs, and phone numbers. Use this MCP to immediately vet user input against known fraud patterns, proxies, or malicious activity. It helps developers build robust sign-up flows and automates risk assessment across your entire platform.

A+ Quality Score 100/100

fraud-detection

risk-management

bot-detection

email-verification

identity-verification



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.

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 — connect your AI agent in under 60 seconds.

IPQualityScore (IPQS) MCP

10 tools available

Cloud-hosted on Vinkius

Running a modern application means handling bad data—fraudulent accounts, bot traffic, suspicious links. This connector gives your AI agent the ability to vet user inputs before they impact your system. You can check an IP address for proxy usage or fraud scores; verify if an email is deliverable and legitimate; or analyze a URL to see if it's linked to malware. These checks happen instantly, letting you block bad traffic at the source. When integrating this into your existing agent workflows through Vinkius, you get immediate access to industry-leading security intelligence. This lets developers build trust and keep user bases clean automatically.

Core Capabilities

01 — Vet User Emails

Checks an email address for fraud risk and determines if it's safe for new account registrations.

03 — Scan URLs

Returns a risk score and classification for any website link provided by users, catching malicious sites.

02 — Analyze IP Addresses

Provides detailed scores on an IP, flagging potential proxies, VPN use, or signs of automated bot activity.

04 — Check Phone Numbers

Analyzes phone numbers to determine line type and overall risk level for identity verification.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/ipqualityscore-ipqs — connect your AI agent in three steps.

- 01 Your AI client sends the data—like an IP address or a user's email—to this MCP.
- 02 The connector performs a real-time lookup against global fraud and reputation databases.
- 03 You receive structured risk scores, flags, and detailed findings that tell you if the input is safe to use.

The bottom line is your AI client gets instant security intelligence on any user-provided data point, letting you make smart decisions immediately.

Built For

Security engineers and product managers who spend time manually reviewing suspicious sign-ups or auditing payment flows. If you're tired of building custom regex checks that fail against modern fraud tactics, this is for you.

Security Engineer

Uses the IPQS MCP to automatically flag high-risk IPs during penetration testing and build automated threat detection rules.

Full-Stack Developer

Integrates email_lookup and ip_lookup into sign-up forms, ensuring that only legitimate users can create accounts on the platform.

Product Manager

Defines new security requirements for user onboarding flows, using this MCP to prove risk reduction before implementation.

What Changes When You Connect

- 01 Stop fraud at the door. When a new user signs up, use email_lookup and ip_lookup to immediately score their credentials against known bot or fraudulent patterns.

-
- 02 Manage high-risk links automatically. Any time a user submits an external URL—maybe in a profile bio or a message—run `url_lookup` before letting them post it.

 - 03 Audit account health with built-in tools. You can call `get_credits` and `list_stats` to keep track of your service usage without needing custom dashboards.

 - 04 Verify identity using phone number data. Use `phone_lookup` on incoming contact data to confirm if the number is a standard line or potentially spoofed.

 - 05 Build compliance reports easily. Calling `list_fraud` and `list_reports` lets your agent gather necessary security logs for quick auditing purposes.
-

Real-World Applications

Onboarding a New Enterprise Client

A developer needs to ensure that all incoming API requests from a new client are legitimate. Instead of manually checking IPs, the agent first calls `ip_lookup` on every source IP and then uses `email_lookup` on all associated contacts, guaranteeing high quality leads before they enter the pipeline.

Preventing Account Takeover

A user attempts to reset their password using an old or suspicious phone number. The agent uses `phone_lookup` on the provided number; if the risk score is too high, it forces the user through manual identity verification instead of proceeding with the password reset.

Moderating User-Generated Content

A content moderator finds a suspicious link posted in a forum thread. The agent intercepts the URL and runs `url_lookup` immediately. If the score is high, the system blocks the post and alerts a human reviewer.

E-commerce Fraud Investigation

An analyst needs to check a batch of conversion data. They use `list_conversions` combined with calling `list_fraud` logs to correlate suspicious purchase patterns with known fraud events, speeding up manual investigations.

Patterns to Avoid

Writing custom blocklists

X AVOID

Trying to maintain a local database of bad IPs or spam emails by copying and pasting findings from security forums. This process is slow, incomplete, and never keeps up with modern fraud tactics.

✓ INSTEAD

Use the IPQS MCP's real-time lookups. Instead of managing blocklists manually, let your agent call `ip_lookup` on every incoming connection or `email_lookup` on every new signup. The service handles the global threat intelligence.

Relying only on basic regex

X AVOID

Using simple code checks to validate an email address format (e.g., 'text@domain.com'). This approach fails instantly when encountering throwaway or temporary email services.

✓ INSTEAD

Run the `email_lookup` tool. It goes beyond formatting and actually analyzes the domain's reputation, deliverability, and historical fraud scores.

Ignoring context

X AVOID

Treating all user inputs equally. A link submitted in a private message gets the same vetting as a core API parameter, risking exposure to malicious payloads.

✓ INSTEAD

The agent must decide when to check and what tool to use. For links, call `url_lookup`; for IPs, call `ip_lookup`. Use the right security context for every piece of data.

The Right Fit

Use this MCP if your primary pain point is validating *reputation*—whether a piece of data (IP, URL, email, phone number) has been seen before in connection with fraud or malicious activity. You need external, real-time threat intelligence to secure user onboarding and content moderation. Don't use it if you simply need to format or validate the *syntax* of the data; for that, a simple regex check works fine. However, if you are building any public-facing system where bad actors could exploit gaps in your security, this MCP is essential. For instance, if you only care about knowing how many times a feature was used, `list_stats` suffices. But if you need to know *who* used it and *if they were fraudulent*, you must use the lookup tools.

Manually vetting user input is a nightmare of clicks and spreadsheets.

Today, every time a new account signs up or a user posts a link, someone has to manually check if that data looks suspicious. You're clicking into one dashboard for IP checks, then copy-pasting the email address into another service just to run validation, and finally pasting the URL into a third security checker. It's slow, it's error-prone, and you're always hours behind the fraudster.

With this MCP, your agent handles all of that in one go. You send the suspicious data point—be it an IP address or an email—and get back a clean, actionable risk score instantly. Your workflow doesn't stop; security checks just become part of the normal flow.

Get instant fraud scores with the IPQualityScore (IPQS) MCP

The tedious parts that vanish are the cross-service data transfers. You don't copy an email from your sign-up form into a separate validation tool; you simply pass it to the agent, which uses the `email_lookup` tool and returns the score right in the conversational response.

What changes is trust. Instead of guessing if user input is safe, you know for sure. You can enforce security checks on every single piece of data coming into your platform.

IPQualityScore (IPQS) MCP: 10 Tools

Analyze IP addresses, email addresses, phone numbers, and URLs instantly to score them against global databases of fraud and risk.

#	TOOL	DESCRIPTION
01	<code>email_lookup</code>	Analyzes an email address and returns a fraud risk score, helping prevent fraudulent accounts during registration.
02	<code>get_account</code>	Retrieves current usage details about your IPQS account plan and configuration status.
03	<code>get_credits</code>	Provides the remaining credit balance, ensuring the service stays active within your allocated quota.
04	<code>ip_lookup</code>	Returns fraud scores, proxy/VPN detection results, and geographical data. Essential for identifying malicious users or automated bots during sign-up or transaction processes. Analyzes an IP address for fraud and proxy detection.
05	<code>list_conversions</code>	Lists records of tracked conversions, useful for auditing e-commerce or affiliate marketing activity.
06	<code>list_fraud</code>	Gathers recent logs detailing fraud events and security triggers for high-level threat monitoring.
07	<code>list_reports</code>	Provides a list of recent security reports and audit findings related to platform usage.
08	<code>list_stats</code>	Lists overall account usage statistics, helping you monitor the health and volume of your integration.
09	<code>phone_lookup</code>	Analyzes a phone number to return its line type and an indicator of potential fraud risk for identity verification.
10	<code>url_lookup</code>	Returns a comprehensive risk score and classification for any URL, useful when auditing suspicious links.

See It in Action

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

U Analyze the IP address 8.8.8.8 for fraud risk.



I'll perform an IP lookup using IPQS to check for any malicious indicators.

U Verify if the email address 'fraudster@test.com' is risky.



I'll use IPQS email validation to check the risk score for that address.

U Check this URL for potential malware: <http://malicious-site.com>



I'll perform a malicious URL lookup with IPQS for you.

Frequently Asked Questions

01 How often do I need to call the IPQualityScore (IPQS) MCP?

You should use this MCP anytime user-provided input is critical, such as during sign-up flows or when processing external links. Calling it proactively prevents fraud before it happens.

02 Does IPQualityScore (IPQS) help with bot detection?

Yes. The ip_lookup tool specifically analyzes IPs for signs of automated activity, detecting if a user is connecting through known proxies or VPNs used by bots.

03 Can I use this MCP for e-commerce auditing?

Absolutely. You can call `list_conversions` and `list_fraud` to audit purchase patterns and identify potential fraudulent transactions on your platform.

04 What is the difference between `ip_lookup` and `url_lookup`?

`ip_lookup` checks the source of a connection (the IP address) for fraud or proxy detection. `url_lookup` analyzes the content of a link itself to score its malicious potential.

05 How do I check if an email is spam using IPQualityScore (IPQS)?







Use the `email_lookup` tool. It doesn't just check for syntax; it analyzes the domain and history against global fraud databases to give you a proper risk score.

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 <code>"ipqualityscore-ipqs": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

IPQualityScore (IPQS) 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 · 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 IPQualityScore (IPQS). 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	IPQualityScore (IPQS) MCP
Server ID	019d75bb-afd0-72c6-9bb6-8a45451c8b55
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/ipqualityscore-ipqs.