MCP SERVER | NO CODE | CLOUD HOSTED

SEON MCP

Automate Fraud and Compliance Checks

SEON MCP connects advanced fraud and risk intelligence directly to your chat interface. It gives you real-time access to digital footprints for emails, phone numbers, and IPs. You can run full Know Your Customer (KYC) checks, screen against global watchlists for AML compliance, and monitor transactions instantly, all without leaving your workflow.

A+ Quality Score 100/100

fraud-prevention | risk-scoring | aml-compliance | device-fingerprinting | identity-verification | data-security

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

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.

SEON MCP

12 tools available
Cloud-hosted on Vinkius

Using this MCP lets you manage fraud prevention and risk orchestration through natural conversation with your AI agent. Instead of juggling multiple dashboards or running manual background checks, your agent handles the heavy lifting. You can ask it to assess a user based on their email address alone, which instantly pulls in digital metadata for social links and domain age. It screens users against global sanctions lists for AML compliance, analyzes IP risk scores, and even lets you monitor organizational fraud health by listing rules or reviewing past transactions. This capability means your AI acts like a dedicated, always-on risk analyst that lives right inside your chat client, making complex operations simple to execute within the Vinkius catalog.

Core Capabilities

01 — Run Fraud Risk Assessments

The MCP performs comprehensive fraud checks on any given transaction or user registration data.

02 — Analyze Digital Identities

It retrieves social media and domain metadata for emails, phone numbers, and IP addresses to build a digital profile.

03 — Verify Compliance Status

The agent screens individuals against global watchlists like sanctions and PEP lists for AML compliance.

04 — Manage Risk Data Sets

You can list existing fraud rules or add specific items to custom blacklists and whitelists.

05 — Track Business Operations

The MCP allows you to get current account details, view transaction specifics, and monitor the status of various AML monitors.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP on Vinkius and provide your SEON Admin API Key from your dashboard settings.

02 Tell your AI agent what you need. For example, ask it to 'check the risk for this email' or 'run fraud assessment on this IP address.'

03 Your agent uses the available tools to run real-time checks and returns a clear, actionable report detailing scores, risks, and compliance flags.

The bottom line is you get complex fraud intelligence delivered conversationally, without leaving your preferred chat application.

Built For

This MCP is for operations teams in high-volume industries like e-commerce and fintech. It's built for the Compliance Officer who is tired of manually cross-referencing sanctions lists, or the Risk Manager who needs instant fraud scores without switching between five different dashboards.

Compliance Officer

Runs automatic AML screening and uses tools like list_aml_monitors to ensure all transactions adhere to global regulatory standards.

Risk Manager

Conducts real-time fraud checks by combining multiple data points, such as checking_email results with check_ip risk scores, for immediate decision making.

Operations Analyst

Manages custom user lists and trains the fraud model by using add_label to feed feedback directly into the system.

What Changes When You Connect

01 Stop jumping between apps. Your AI agent handles everything from checking the phone number (check_phone) to running a full fraud assessment (check_fraud), all in one chat window.

02 Ensure global compliance instantly. Use aml_screening to screen users against sanctions and PEP lists, giving you clear answers on regulatory risk without manual spreadsheet work.

03 Build better models faster. Instead of guessing, use add_label to feed specific feedback directly into the system, ensuring your fraud detection tools train on real-world outcomes.

04 Get comprehensive identity context immediately. When a user signs up, run check_email and get instant social media metadata alongside domain age, improving risk scoring accuracy.

05 Understand complex transactions at a glance. Use get_transaction to pull all relevant details—like card country mismatches or velocity abuse flags—without leaving the chat interface.

Real-World Applications

Investigating suspicious new sign-ups

A risk analyst encounters a high volume of registrations. They ask their agent to check_fraud on all pending accounts and then use check_ip for any IPs with known VPN usage, instantly flagging potential bot activity.

Handling international payment disputes

An ops team member receives a payment from a new country. They ask the agent to get_transaction details to verify shipping/billing matches and run aml_screening on the sender's name, mitigating cross-border financial risk.

Updating internal fraud policies

A compliance officer needs to update rules. They use list_rules to review existing logic, then use add_to_list to block a known fraudulent domain on the blacklist, ensuring all agents respect the change.

Investigating account takeovers

A security team member notices suspicious activity. They check_account_info for user history, then use check_phone and check_email to see if linked profiles suddenly changed or match known compromised data patterns.

Patterns to Avoid

Doing manual background checks

✗ AVOID

Copying a phone number into one dashboard, pasting an email into another, and opening a third tab to check IP risk. This takes minutes and is error-prone.

✓ INSTEAD

Just ask your agent to combine these steps: 'check_fraud for this user, using their email, phone, and IP.' The MCP handles the multi-step data collection automatically.

Over-relying on single signals

✗ AVOID

Seeing a low risk score and approving a transaction without checking if the card country matches the billing address. This opens you up to fraud.

✓ INSTEAD

Always combine checks: Use get_transaction first, then check_fraud with the results. The system aggregates multiple data points for a safer decision.

Forgetting compliance steps

✗ AVOID

Approving an account without verifying the user's identity against global watchlists because it seems like a simple domestic transaction.

✓ INSTEAD

Make aml_screening a mandatory step for every new high-value user. It prevents regulatory fines and keeps you compliant.

The Right Fit

Use this MCP if your risk workflow requires constant, multi-layered data verification across identity, transactions, and compliance domains. If your core problem is that fraud detection involves checking five different types of identifiers (email, phone, IP, account history, watchlist status), you need this tool. Don't use it if you only need to list simple users or retrieve basic contact information; those are single-purpose tools. Instead, use a dedicated identity management MCP for isolated tasks like listing users or retrieving only the email footprint via check_email.

This MCP shines when your decision depends on combining multiple pieces of data—for instance, 'Is this IP high risk *and* is the user's phone number linked to a known watchlist?' The ability to orchestrate these complex questions in natural language is its biggest advantage.

The Daily Grind: Manual Risk Assessment

Right now, checking a single high-risk user requires clicking through at least three different systems. You copy the email into the identity tool, open another tab to run an IP check, and then switch back to the compliance dashboard just to run aml_screening against sanctions lists. It's slow, it's tedious, and you inevitably lose context or miss a critical cross-reference.

With this MCP, your agent does all that work for you. You simply tell it what needs screening—for example, 'Assess the risk of this user.' The tool runs multiple checks in sequence, pulling together digital footprint data and compliance scores into one clear report.

Get Full Control with SEON

You no longer have to manually manage fraud rules or wait for a developer to update the whitelists. You can use list_rules to see exactly what's running, and then immediately use

What changes is that you are taking full control of your risk data in real-time. Your AI agent acts as an instant compliance officer, making complex governance decisions conversational.

add_to_list to enforce a new block on a suspicious IP.

SEON: 12 Risk Intelligence Tools

These tools allow you to perform detailed operations like aml_screening, check_fraud, and list_rules directly through natural conversation.

#	TOOL	DESCRIPTION
01	`add_label`	Adds a specific label or classification tag to existing data for training purposes.
02	`add_to_list`	Adds an item, like a phone number or email, to either a custom blacklist or whitelist.
03	`aml_screening`	Performs comprehensive Anti-Money Laundering screening against global watchlists.
04	`check_email`	Gets a detailed digital footprint and risk score for any given email address.
05	`check_fraud`	Runs a full-scope fraud check on a transaction or user registration attempt.
06	`check_ip`	Retrieves risk data and location intelligence associated with an IP address.
07	`check_phone`	Gathers intelligence on a phone number, including carrier details and validity checks.
08	`get_account_info`	Retrieves general account profile information for monitoring purposes.
09	`get_transaction`	Fetches specific, detailed metadata about a completed or pending transaction.
10	`list_aml_monitors`	Displays the list of current Anti-Money Laundering monitoring rules in place.
11	`list_lists`	Shows all custom blacklists and whitelists currently managed by your account.
12	`list_rules`	Lists all active fraud rules that govern risk scoring for the platform.

See It in Action

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

U Check this email for fraud risk: suspicious.user@example.com.

I've retrieved the digital footprint. The email has no linked social media accounts and the domain was registered 2 days ago, resulting in a high risk score of 85. Would you like to check the associated IP address?

U Run a fraud risk assessment on a new user registration with email and IP analysis.

Fraud assessment completed. Overall risk score: 23/100 (Low Risk). Email analysis: user@gmail.com - Registered 4 years ago, active social profiles (LinkedIn, GitHub), data breach count: 1 (minor). Email score: 15/100 (safe). IP analysis: 203.0.113.45 - ISP: Comcast, Location: San Francisco, CA. VPN: No. Tor: No. Proxy: No. IP score: 18/100 (safe). Device fingerprint: Chrome 125, macOS, unique device. Phone: +1-555-0198 - Carrier: AT&T, type: mobile, valid format. Recommendation: APPROVE.

U Show me all flagged transactions from the last 24 hours with their risk scores and reasons.

12 flagged transactions in 24 hours. High risk (3): TX-8923 ($2,400, score 89, disposable email + VPN), TX-8934 ($1,800, score 82, velocity abuse - 5 orders in 10 min), TX-8941 ($3,200, score 91, card-country mismatch + new account). Medium risk (5): velocity patterns, mismatched billing/shipping (scores 55-72). Low risk (4): minor flags, first-time buyers (scores 35-48). Auto-blocked: 2 transactions. Manual review queue: 6. False positive rate (last 30 days): 3.2%.

Frequently Asked Questions

01 **How do I use SEON MCP to check a user's email for fraud?**

You simply instruct your agent to perform a check_email on the address. The tool will return a complete digital footprint, showing things like linked social profiles and domain age alongside an immediate risk score.

02 **What is aml_screening doing when I use the SEON MCP?**

aml_screening runs your user against global sanction lists and watchlists. It confirms if the individual or entity poses a regulatory compliance threat, which is vital for financial institutions.

03 **Can I update my fraud models with this MCP?**

Yes. You use add_label to give feedback on past transactions. This trains your underlying models and helps the platform stay synchronized with emerging fraud patterns.

04 **What if I need to check an IP address for risk?**

Use check_ip. It retrieves detailed data points, including the ISP, geographic location, and whether the IP is known to be associated with VPNs or proxies.

05 **Does SEON MCP only work for e-commerce fraud?**

No. While excellent for transaction risk, it also helps compliance officers by allowing them to list_aml_monitors and track general account details using get_account_info.

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 `"seon": { "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

SEON 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 SEON. 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	SEON MCP
Server ID	`019dd158-1ed5-733d-9769-4bed0bfc5955`
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/seon.