MCP SERVER | NO CODE | CLOUD HOSTED

DataDome MCP for AI Agents

Govern web and API traffic by monitoring bot threats and fraud activity

DataDome connects your AI client directly to enterprise-grade bot protection and fraud prevention data. It lets you audit protected applications, list recent bot attacks, and get real-time statistics on suspicious traffic without manually logging into a dashboard.

A+ Quality Score 100/100

bot-protection | fraud-prevention | threat-intelligence | web-security | endpoint-protection | anti-scraping

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

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.

DataDome MCP

10 tools available
Cloud-hosted on Vinkius

Your web and mobile apps are under constant automated attack from scrapers and bots. Instead of navigating complex security dashboards to figure out what happened last night, this MCP lets your AI client talk directly to DataDome's API. It pulls together everything you need: it monitors which endpoints are struggling, tracks the types of malicious traffic hitting your site, or retrieves a quick count of allowed vs. blocked requests today. This allows you to understand protection status and threat trends right inside your workflow. Through Vinkius, you get that centralized control without having to switch context. You can use natural language commands to check detailed logs on specific attacks or manage custom security rules.

Core Capabilities

01 — Summarize bot traffic patterns

Get a categorized breakdown of all incoming traffic, separating legitimate search engine bots from malicious scrapers.

02 — Check endpoint health and performance

Verify the current operational state of any protected web or mobile endpoint, including latency rates and error counts.

03 — Retrieve real-time protection statistics

Access live data on allowed requests versus blocked attempts, captcha pass rates, and identified bot categories.

04 — Get detailed threat forensics

Resolve full technical details for a specific security incident, including request headers and behavioral patterns.

05 — List all protected applications and endpoints

View every application (web or mobile) integrated with DataDome, as well as metadata for every secured endpoint URL.

06 — Review recent access logs and incidents

Generate a stream of historical requests processed by the system, showing bot scores, decisions, and geo-location data.

07 — Manage custom security rules

View all currently active custom detection rules, including which IPs or User-Agents are allowed or blocked.

One Click on Vinkius — From Prompt to Execution

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

01 Connect this MCP to your AI client and authorize it using your DataDome Management API Key.

02 Your agent then uses natural language prompts to request specific data, such as the status of a key endpoint or logs from yesterday.

03 The system executes the necessary API calls, pulling structured security metrics directly into your conversation.

The bottom line is that you talk to DataDome's protection layer using plain English commands instead of navigating dashboards and writing CURL requests.

Built For

This MCP is for security operations teams who live in a state of constant vigilance. If your job involves monitoring attack surfaces, tracking fraud vectors, or validating application health across multiple endpoints, this tool cuts out the manual dashboard clicking and reporting.

Security Engineer

Checks threat logs instantly to verify endpoint protection status or investigates specific suspicious activity using detailed request headers.

SRE / DevOps Specialist

Monitors application health and bot traffic trends by querying the system for real-time latency metrics or overall traffic summaries.

Fraud Analyst

Gathers structured data on recent bot activities, including threat types like credential stuffing, to build detailed incident reports.

What Changes When You Connect

01 Stop manually cross-referencing dashboards. Use list_recent_threats to pull a clear, immediate summary of the latest suspicious activities directly into your chat.

02 Track performance without guesswork. Call get_endpoint_health on any protected URL to check latency and error rates in real time, ensuring uninterrupted service.

03 Understand attack volume instantly. Retrieve protection statistics using get_protection_stats to see live counts of blocked requests versus allowed traffic for immediate capacity planning.

04 Deep dive into attacks with pinpoint accuracy. Use get_threat_details when an incident occurs; you'll get the full technical breakdown, including request headers and detection logic.

05 Manage your defense strategy from chat. List custom bot rules lets you quickly audit who has allowed or blocked specific IPs or User-Agents without logging into the management console.

Real-World Applications

Investigating a sudden spike in errors

An SRE notices the error rate jumped yesterday. They ask their agent to run get_endpoint_health on the main API gateway. The agent immediately reports that while latency is stable, the error count spiked because of an unlisted mobile endpoint.

Forensically analyzing a suspicious login attempt

A fraud analyst spots unusual activity and asks to see recent access logs filtered for that IP. The agent uses list_access_logs, showing the bot score was high and detailing the failed credential stuffing attempts.

Auditing anti-scraping rules

A security engineer needs to verify if a competitor successfully bypassed their defenses. They run list_custom_bot_rules, confirming that the 'crawler' rule is still active and listing the criteria used.

Getting traffic context for a meeting

A product manager needs to present data on bot activity. Instead of compiling reports, they ask for get_bot_traffic_summary, which provides an immediate, categorized breakdown showing 85% search engine bots and 15% scrapers.

Patterns to Avoid

Treating the MCP like a database lookup

✗ AVOID

Asking the agent for 'all data from last month' without specifying what kind of activity (e.g., list_access_logs). This results in an unusable, massive dump of raw log entries.

✓ INSTEAD

Always narrow your focus. Instead, ask to list recent threats using list_recent_threats or search for specific attack types with search_threats_by_type. Focus on the *outcome*, not just the data source.

Trying to configure rules via chat

✗ AVOID

Telling the agent, 'Block all traffic from IP 123.' The MCP can only read and report data; it cannot modify configurations or write new rules.

✓ INSTEAD

Use list_custom_bot_rules first to see existing rules. If you need to change a rule, you must do so in the DataDome dashboard, but your agent can confirm what the current configuration is.

Ignoring endpoint health metrics

✗ AVOID

Assuming an endpoint is fine because it's online, without checking performance. This overlooks high latency or sudden spikes in error rates.

✓ INSTEAD

Always run get_endpoint_health when reviewing a critical service. This provides objective measurements of latency and operational status that simple uptime checks miss.

The Right Fit

Use this MCP if your primary pain point is the time spent correlating security metrics across different dashboards. If you need to know *what* happened, *when*, and *why*—like investigating a specific threat or auditing a rule set—this MCP is perfect. However, don't rely on it for actual configuration changes; this tool reads data, but it doesn't write rules. If your goal is to actually change the blocking behavior (e.g., creating a new custom bot rule), you still need to use DataDome's primary management interface. Use list_protected_endpoints first if you don't know which services are even monitored.

DataDome MCP: Monitoring Web and API Fraud Traffic

Today, fraud analysts spend hours logging into multiple security dashboards. They copy threat IDs from one tab, cross-reference them with access logs in a second dashboard to find the origin IP, and then manually check the endpoint health status of that service. It's tedious clicking through tabs just to build a single incident report.

With this MCP, you simply ask your agent about the incident. You can request threat details for an ID or list recent threats by type, getting all the necessary forensic data points—the IP, the behavior pattern, and the timestamp—in one conversation. It cuts out the entire manual investigation loop.

DataDome MCP: Auditing Protection Status and Rules

Before a major site release or migration, security teams must verify that every single endpoint is covered. Manually checking the protection status for dozens of URLs across web and mobile APIs is a massive chore prone to human error.

Now, you can list all protected endpoints in one go, getting an immediate overview of which services are secured and what type of protection they have enabled. This gives you instant coverage confirmation that used to take hours of spreadsheet auditing.

DataDome: 10 Tools for Analyzing Bot Traffic and Endpoint Metrics

Use these tools to get immediate metrics on bot activity, check endpoint health, or retrieve detailed protection statistics through natural language queries.

#	TOOL	DESCRIPTION
01	`get_bot_traffic_summary`	Returns a summary of traffic from 'Good Bots' versus 'Bad Bots', showing their combined impact on total site traffic.
02	`get_endpoint_health`	Checks the health status of a specific protected endpoint, providing latency metrics and error rates.
03	`get_protection_stats`	Retrieves real-time protection statistics, including counts for allowed requests, blocked attempts, and captcha pass rates.
04	`get_threat_details`	Gets full technical details for a specific threat ID, revealing the behavioral patterns and detection logic that triggered it.
05	`list_protected_applications`	Lists all applications (mobile or web) integrated with DataDome along with their protection types.
06	`list_protected_endpoints`	Returns metadata for every endpoint protected by DataDome, including the URL and its current status.
07	`list_access_logs`	Lists recent access logs processed by DataDome, filtered to show bot scores, decision outcomes, and geo-location data.
08	`list_custom_bot_rules`	Returns a list of custom detection rules configured in your account, including match criteria and action type.
09	`list_recent_threats`	Lists recent security incidents detected by DataDome, detailing the threat types, origin IPs, and timestamps.
10	`search_threats_by_type`	Searches through recent incidents to isolate specific attack vectors based on a keyword like 'scraper' or 'crawler'.

See It in Action

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

U What were the top 3 types of bot threats detected in the last week?

Threat Summary: Last 7 Days

- ❌ **Scraping Activity:** 4,102 incidents (Primary concern)
- 🤖 **Credential Stuffing:** 985 attempts (Focus on high-value targets)
- 🐌 **Slow Crawling:** 610 events (Origin: Search Engine Bots)

Action: Would you like to see the top originating IPs for the scraping activity?

U Check the health and status of our main checkout API endpoint.

Endpoint Status Report: CheckoutAPI

METRIC	VALUE
Status	✅ Healthy
Latency (Avg)	42ms
Error Rate	0.1%

Note: Protection is active, and the last recorded spike was minor. The current bot traffic summary shows a slight increase in 'Scraper' activity today.

U Give me a summary of all suspicious requests from yesterday afternoon.

Incident Log Summary: Yesterday (12 PM - 5 PM)

- **Total Requests Processed:** 1.8 Million
- **Blocked by DataDome:** 34,500
- **Top Threat Type:** Scraper (70% of blocked activity)
- **Highest Risk IP Blocked:** `203.0.113.4` (Type: Credential Stuffing).

You can ask me to get the full technical details for that specific block if you want more info.

Frequently Asked Questions

01 **How does DataDome MCP help audit my protected web endpoints?**

It lets your agent query all your secured URLs to check their current health, latency, and protection status. You'll get real-time metrics on performance without having to manually log into the dashboard.

02 **Can I use DataDome MCP to analyze bot traffic patterns?**

Yes, you can ask it for a summary of all incoming traffic. It breaks down bots by type—like search engines versus malicious scrapers—giving you an instant picture of your threat landscape.

03 **What kind of security incidents can DataDome MCP report on?**

You can list recent threats, getting details like the originating IP, what type of attack occurred (e.g., scraping), and exactly when it was detected. This is vital for incident reporting.

04 **Is DataDome MCP useful for tracking API protection metrics?**

Absolutely. You can retrieve real-time statistics on your APIs, getting counts of blocked requests versus allowed ones, helping you understand the volume and type of activity hitting your back end.

05 **How do I check if my custom bot rules are working correctly with DataDome MCP?**

You can use list_custom_bot_rules to see all active rules. You can then cross-reference those rules with recent access logs to verify that the correct actions (allow/block) were taken for specific traffic types.

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 `"datadome": { "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

DataDome 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 DataDome. 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	DataDome MCP
Server ID	`019d7581-f3aa-7087-88d3-9317b50dbe97`
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/datadome.