MCP SERVER | NO CODE | CLOUD HOSTED

Trend Micro MCP

Correlate alerts with endpoint activity using natural language.

Trend Micro MCP lets your AI client investigate security threats directly from your Vision One infrastructure. Instead of navigating complex SIEM dashboards or writing custom API scripts, you talk to it naturally. It gives you immediate access to high-fidelity telemetry, XDR detections, and structural alerts. You can check suspicious URLs, list all deployed endpoints, and hunt forensic logs—all through plain language conversation.

A+ Quality Score 100/100

cybersecurity | threat-intelligence | xdr | endpoint-security | network-security | vulnerability-scanning

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/trend-micro

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.

Trend Micro MCP

8 tools available
Cloud-hosted on Vinkius

Connect your AI agent directly into your Trend Micro Vision One security system. This MCP lets analysts bypass clunky dashboards and complicated interfaces, allowing them to interact with raw threat data using only natural language. You don't need to know the API structure or spend time writing scripts just to get basic intel.

Need to understand a potential breach? Ask your agent for details on a specific alert ID. Want to see what machines are connected to the network? Just ask it to list all managed endpoints. Your agent can pull forensic logs around targeted emails, check live indicators of compromise like suspicious IPs or URLs, and even review raw detections that haven't triggered an official alert yet.

This capability lets your Security Operations Center (SOC) team move faster when responding to incidents. It’s the kind of focused power you only get by connecting through a central hub like Vinkius, giving your agent instant access to thousands of security tools and data sources.

Core Capabilities

01 — List current structural alerts

It pulls an immediate list of all active security alerts from the Trend Micro Vision One workbench.

02 — Review specific alert details

You can drill down into a single, problematic alert ID to see exactly what triggered it and evaluate its potential impact.

03 — Check network assets

The agent lists all physical devices that are deployed and managed within your organization's network sphere.

04 — Identify threat indicators

It queries live data to show any suspicious objects, such as blacklisted URLs, malicious IP addresses, or file hashes found in your network.

05 — Search deep activity logs

You can instruct the agent to hunt through detailed endpoint processes or specific email workflow histories for forensic evidence.

One Click on Vinkius — From Prompt to Execution

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

01 First, activate this MCP connector within your organization's security workspace.

02 Next, provide a secure API Key generated inside the Vision One console, along with your specific AWS or Cloud region code.

03 Finally, engage your AI agent and ask it for an immediate status check on your domain's health.

The bottom line is you get to analyze complex security data using a simple conversation instead of complicated dashboards.

Built For

This MCP is built for the technical experts who spend their day staring at dozens of tabs and wrestling with API documentation. It helps the SOC Analyst tired of clicking through endless SIEM records, or the Threat Hunter who needs to correlate an IP address with endpoint activity in seconds.

Security Operations Center (SOC) Analyst

They use this MCP during incident response to gather associated observables and forensic logs through rapid conversation, accelerating their ability to contain a threat.

Threat Hunter

They query the system instantly for specific indicators of compromise, like untrusted blacklisted URLs, without manually running multiple searches across different tools.

Security IT Engineer

They validate whether a newly deployed endpoint was accurately tracked and successfully integrated just by asking the agent to check asset status via terminal command.

What Changes When You Connect

01 Stop manually navigating dashboards. Instead of clicking through five different tabs to get a full picture, you ask your agent to list all active structural security alerts and immediately understand the scope.

02 Speed up forensic analysis. If something suspicious happens, asking for detailed endpoint activity logs lets you trace exactly what processes ran on the device without writing complex query language.

03 Get comprehensive visibility into assets. You can quickly use list_managed_endpoints to verify if a new machine has been successfully tracked and integrated into your security monitoring.

04 Improve threat intelligence depth. Rather than guessing, you can use list_suspicious_objects to check live indicators of compromise for URLs or IPs against known blacklists.

05 Simplify investigation scope. Your agent groups related data points, allowing you to jump straight from a general alert ID (using get_alert_details) to the underlying network observables that matter.

Real-World Applications

Investigating an Alert Spike

A SOC analyst sees a high-severity alert. Instead of opening five different consoles, they simply ask their agent for details on the specific alert ID and then follow up by running list_endpoint_activity_logs to see what happened right before the alert fired.

Validating New Assets

A security engineer needs proof that a newly deployed laptop is fully covered. They run list_managed_endpoints and check the output to confirm the asset's status, ensuring it’s tracked correctly in Vision One.

Tracking Phishing Campaigns

A threat hunter suspects lateral movement via email. They ask for logs on email activity (list_email_activity_logs) and then use list_suspicious_objects to check if the malicious URLs mentioned in the emails are already known bad IPs or domains.

Deep Dive Forensics

A user needs to understand a breach. They ask their agent to look at raw detections (list_recent_detections) and then request list_endpoint_activity_logs for the machine involved, getting a clean timeline without sifting through massive JSON files.

Patterns to Avoid

Writing complex API calls

✗ AVOID

I have to write an endpoint that queries alert IDs AND then another one that gets the associated logs, and I have to map them myself.

✓ INSTEAD

Just ask your agent for 'details on active alerts related to machine X.' It handles listing security alerts and getting the specific details (get_alert_details) in a single conversation.

Switching between dashboards

✗ AVOID

I need to check suspicious IPs, but I have to leave the alert dashboard and go into the threat intelligence panel just for that one piece of data.

✓ INSTEAD

Ask your agent to list_suspicious_objects. It checks both the threat intel feed and reports it back right where you are.

Assuming endpoint status

✗ AVOID

I think this machine is online, but I can't remember if the last time I checked was accurate.

✓ INSTEAD

Use list_managed_endpoints. It gives a definitive, up-to-date roster of every connected asset and its current health status.

The Right Fit

You should use this MCP if your team's security process involves correlating multiple data sources—like linking an alert to an endpoint process log or checking an IP against a threat feed. It excels when you need deep, forensic visibility without writing code. Don't use it if all you need is a simple daily dashboard summary; for basic monitoring, a standard SIEM view might be fine. However, if your job requires correlating list_managed_endpoints data with recent

detections or pulling specific alert details (get_alert_details) to understand the 'why', this MCP is essential. It's designed for active investigation, not passive viewing.

The constant pivot between dashboards and consoles sucks.

Right now, investigating a single alert means logging into five different panels: the alerts dashboard, the asset inventory panel, the threat feed, the log viewer, and then maybe an email system. You click through tabs, copy unique IDs from one screen, paste them into another, run a query, wait for it to load, and then finally piece together what happened.

With this MCP, you just talk to your agent. You can ask about 'all alerts related to suspicious IPs' and the system pulls the data—alert details, asset status, and threat intelligence—and gives you one clean answer. It handles all that cross-panel correlation automatically.

Trend Micro MCP: Correlate alerts with endpoint activity using natural language.

The manual process of hunting for suspicious URLs often means running one query on the threat feed, then checking a separate dashboard for recent detections to see if that URL was hit. It's slow, and you might miss connections because you had to run two different reports.

Now, ask your agent to list_suspicious_objects and cross-reference those findings with the logs from list_email_activity_logs. You get a single, comprehensive report showing exactly where that malicious object landed in your environment.

Trend Micro MCP with 8 Tools

These tools allow your AI client to perform specific, deep-dive actions across Trend Micro Vision One, covering everything from asset tracking to threat intelligence checks.

#	TOOL	DESCRIPTION
01	`get_vision_one_account`	Retrieves your Trend Micro account name and checks the overall connectivity status.
02	`get_alert_details`	Fetches specific details for any single security alert identified by its unique ID.
03	`list_security_alerts`	Generates a list of all current structural alerts recorded in the Vision One workbench.
04	`list_recent_detections`	Pulls a feed of recent security detections from XDR, even if they haven't been promoted to an active alert.
05	`list_email_activity_logs`	Searches the logs for detailed activity related to email workflows, useful for threat hunting.
06	`list_endpoint_activity_logs`	Retrieves telemetry and logs showing what has happened on a specific endpoint device.
07	`list_managed_endpoints`	Lists all the physical assets and devices that are connected to and managed by Vision One.
08	`list_suspicious_objects`	Checks the threat intelligence database for suspicious network objects like IPs, URLs, or files.

See It in Action

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

U Check and list my managed endpoints connected to Vision One right now.

I retrieved 4 endpoints connected successfully. All sensors mark online and active status, running Windows 11 Enterprise nodes.

U Extract the details of the active security alert tagged with ID 22b-88cx.

Retrieved alert data 22b-88cx. Classification: HIGH SEVERITY. Summary implies potential lateral movement via RDP protocol tied to suspicious endpoint 'LAPTOP-HR-04'. Investigate immediately.

Frequently Asked Questions

01 **How do I check my assets using Trend Micro MCP?**

You use the list_managed_endpoints tool to generate an accurate roster of all connected physical and virtual devices. This confirms which machines are currently visible and monitored by Vision One.

02 **Can Trend Micro MCP tell me about suspicious IPs?**

Yes, you ask the agent to list_suspicious_objects. It queries your threat intelligence feed for any blacklisted or compromised IP addresses found within your network's observed traffic.

03 **What is the difference between list_recent_detections and list_security_alerts?**

List_security_alerts focuses only on events that have been formally classified as high-severity alerts. List_recent_detections shows a broader feed of all detections, including low-level activities that haven't reached alert status yet.

04 How do I find logs for an old security incident with Trend Micro MCP?

You can use list_endpoint_activity_logs to search the telemetry data. This allows you to pull specific process details or actions that occurred on a device at a precise time, even if no alert was triggered.

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 `"trend-micro": { "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

Trend Micro 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 Trend Micro. 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	Trend Micro MCP
Server ID	`019d7615-ae2f-732e-8090-313558504fdc`
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/trend-micro.