MCP SERVER | NO CODE | CLOUD HOSTED

Glassnode On-chain Data MCP

Analyze Bitcoin and Ethereum metrics instantly.

Glassnode (On-chain Data) gives your AI client direct access to institutional-grade on-chain metrics for Bitcoin, Ethereum, and over a thousand assets. You can pull real-time exchange flows, network health data, and historical time-series data using natural conversation. It's built for serious crypto research, allowing you to backtest models and compare market trends across multiple digital assets instantly.

A+ Quality Score 98.33/100

on-chain-data | market-intelligence | crypto-analytics | financial-metrics | institutional-data | network-health

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/glassnode-on-chain-data

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.

Glassnode (On-chain Data) MCP

6 tools available
Cloud-hosted on Vinkius

Use this MCP to analyze crypto markets with the precision of institutional tools. Instead of checking ten different dashboards just to track a few key metrics, your AI client fetches everything—from active address counts to specific asset flows—all in one place. You can query thousands of metric paths and pull historical data for any number of assets at once. It's about getting clean, structured time-series data directly into your chat environment. If you use Vinkius, this MCP integrates that depth of on-chain intelligence right where you're working with your AI agent. This lets crypto traders monitor whale movements without ever leaving the conversational interface, and it gives data scientists the raw metrics they need for rigorous modeling.

Core Capabilities

01 — List supported assets

It shows you every asset and blockchain type that can be analyzed.

02 — Explore metric documentation

You get detailed explanations, allowed parameters, and descriptions for specific metrics paths.

03 — Retrieve time-series data

It fetches historical trends for any specified on-chain metric over a chosen period.

04 — Compare multiple assets in bulk

You can pull the same type of metric across several different digital assets simultaneously to compare market health.

05 — Access historical snapshots

It pulls immutable data points for backtesting, ensuring your analysis avoids looking ahead at future prices.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/glassnode-on-chain-data** — connect your AI agent in three steps.

01 Subscribe to this MCP and enter your Glassnode API Key.

02 Direct your AI client to the MCP. Your agent is now ready to receive on-chain data requests.

03 Ask a question in plain language, like 'What were the active addresses for ETH last month?' and get structured metrics returned.

The bottom line is you don't need to manually navigate complex dashboards; your AI client does it for you via conversation.

Built For

Anyone who needs raw, verifiable data about digital asset performance. This is built for the quantitative researcher tired of switching between multiple web portals just to compile a single report.

Quantitative Analyst

They use it to run backtests by fetching point-in-time metrics and calculating complex ratios across various assets.

Crypto Trader

They monitor exchange inflows, tracking large 'whale' movements in real time without leaving their chat interface.

Financial Data Scientist

They pull clean, structured time-series data directly into an analysis environment to build predictive models on network growth.

What Changes When You Connect

01 Precision Backtesting: Use get_pit_metric to pull historical snapshots, eliminating look-ahead bias when testing trading strategies. You're working with accurate data points, not averages.

02 Massive Data Comparison: Forget running separate queries for every asset. The get_bulk_metric tool lets you compare market trends across dozens of assets simultaneously in one go.

03 Instant Documentation Lookup: Need to know what a metric path means? Use get_metric_details to instantly pull parameter rules and descriptions, saving time searching documentation.

04 Full Asset Coverage: Start by running list_assets to see exactly which assets are available. This ensures you never miss a potential data stream for your analysis.

05 Robust Time-Series Analysis: The get_metric tool pulls clean historical data—like active address counts or exchange balances—that's ready to drop straight into a chart or spreadsheet.

Real-World Applications

Tracking Whale Activity

A trader needs to know if large holders are accumulating BTC. They ask their agent, 'Show me the latest exchange flows for BTC.' The agent uses get_metric and returns a clean time-series graph of inflows/outflows.

Comparing Market Health

A financial analyst wants to compare Ethereum's active addresses against Polygon's. They use the agent and ask it to fetch data via get_bulk_metric, getting a side-by-side metric comparison instantly.

Historical Modeling

A data scientist needs to test a strategy from Q4 2023. They use get_pit_metric to lock in the exact state of the market for that date, guaranteeing their model runs on immutable data.

Understanding Metrics

A new user doesn't know if 'market/price_usd_close' is correct. They ask the agent to use get_metric_details and receive a detailed explanation of the metric's parameters and allowed intervals.

Patterns to Avoid

Confusing live data with backtesting

✗ AVOID

Trying to analyze a trading strategy using the most recent, current metrics because it's easier. This introduces look-ahead bias and gives false signals.

✓ INSTEAD

Always use get_pit_metric when simulating past performance. This tool captures an exact snapshot of the market at a specific date, ensuring your backtest is accurate.

Forgetting metric names

✗ AVOID

Spending 20 minutes searching through documentation to find the exact path for 'daily active addresses.'

✓ INSTEAD

First, run list_metrics. This gives you a comprehensive catalog of every possible data point available on Glassnode.

Manual comparison hell

✗ AVOID

Running five separate queries—one for BTC, one for ETH, one for SOL, etc.—and then manually compiling the results into a spreadsheet.

✓ INSTEAD

Use get_bulk_metric. You can ask your agent to gather and present metrics for multiple assets in a single request.

The Right Fit

You should use this MCP if your core need is programmatic, verifiable market data. If you're analyzing relationships between assets (e.g., 'How did ETH active addresses correlate with BTC exchange flows?'), or if you are building a quantitative model that requires time-series depth and historical accuracy, this tool is essential. Don't use it just because you want general news or qualitative commentary; the MCP cannot tell you *why* prices moved, only *what* they were. If your goal is simply to get a high-level summary report for a client meeting, stick with simple web dashboards. But if that report needs underlying metrics—like active addresses (get_metric) or historical price points (get_pit_metric)—this MCP provides the necessary raw materials.

The constant struggle of building crypto reports

Today, getting a comprehensive view of market health is a mess. You're forced to open ten different browser tabs—one for Bitcoin flow, one for Ethereum metrics, another for exchange balances, and then maybe two more just to find the correct historical data point. Then you spend an hour copy-pasting numbers from dashboard A into your Excel sheet, manually verifying that they match what dashboard B says.

With this MCP, you talk to your agent once. You ask it for a multi-asset comparison across three metrics—like active addresses and exchange balances. The raw, structured data arrives formatted and ready in the chat window. It's not just faster; it changes what's possible.

Get On-Chain Metrics with Glassnode (On-chain Data)

The process of looking up metrics used to involve jumping through hoops: checking the API documentation, figuring out parameter limits, and writing complex queries just to list available assets. You'd spend half your time on setup and the other half analyzing.

Now, you simply ask. The MCP handles the complexity behind the scenes. It gives you access to institutional-grade data points like exchange flows and network metrics without needing a single line of code or deep API knowledge.

Glassnode (On-chain Data) MCP: 6 Tools

Use these tools to list supported assets, explore metric documentation, pull time-series data, run bulk comparisons, and access historical snapshots for rigorous crypto analysis.

#	TOOL	DESCRIPTION
01	`get_bulk_metric`	Retrieves a metric across multiple assets at once for comparison.
02	`get_metric_details`	Provides the specific parameters and description for any given on-chain metric path.
03	`get_metric`	Gets a standard time-series data set for one specified metric over time.
04	`get_pit_metric`	Retrieves an immutable historical snapshot of a metric at a specific point in time.
05	`list_assets`	Outputs a comprehensive list of all supported cryptocurrencies and assets on Glassnode.
06	`list_metrics`	Lists every available metric path, allowing you to see what data is even measurable.

See It in Action

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

U List all supported assets on Glassnode.

I've retrieved the list of supported assets. Glassnode currently supports over 1,000 assets including BTC, ETH, and various ERC-20 tokens across multiple blockchains like Bitcoin, Ethereum, and Polygon.

U Get the 'addresses/active_count' metric for BTC from the last 7 days with a 24h interval.

Fetching active address count for Bitcoin... Over the last 7 days, the daily active addresses have fluctuated between 850k and 920k, showing a slight upward trend in network activity.

U Show me the details and allowed parameters for the metric path 'market/price_usd_close'.

The 'market/price_usd_close' metric provides the daily closing price in USD. It supports intervals from 10m to 1month and is available for almost all listed assets. You can filter by 'since' and 'until' timestamps.

Frequently Asked Questions

01 **How do I start getting on-chain data using Glassnode (On-chain Data) MCP?**

First, you need to subscribe and provide your API key. Once connected, simply ask your AI client for the metrics you want; it handles the rest.

02 **Can I use get_bulk_metric if I only care about one asset?**

Yes, you can. While designed for comparison, you can still use it to query a single asset, but listing metrics or using get_metric is often simpler for singular focus.

03 **What is the difference between get_metric and get_pit_metric?**

get_metric pulls standard time-series data (like averages over a period). get_pit_metric provides an immutable snapshot of the market at one precise moment, which is critical for backtesting.

04 **Does list_assets show all assets I can analyze?**

Yes, running list_assets will give you a complete catalog of supported digital assets and blockchains right out of the gate.

05 **What if I need documentation for a metric path? Should I use get_metric_details?**

Yes. Use get_metric_details to understand exactly what a specific metric means, what parameters it accepts, and how it's calculated.

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 `"glassnode-on-chain-data": { "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

Glassnode (On-chain Data) 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 Glassnode (On-chain Data). 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	Glassnode (On-chain Data) MCP
Server ID	`019e389f-fff5-7004-a5a5-af0fbabaac7d`
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/glassnode-on-chain-data.