

M C P S E R V E R

N O C O D E

C L O U D H O S T E D

NYC Open Data Explorer MCP

Query thousands of New York City public datasets directly through your AI.

NYC Open Data Explorer MCP gives your AI client direct access to the city's entire public data catalog. You can search through 3,000+ datasets, inspect specific column metadata, and run filtered queries using SOQL. It works with any MCP-compatible client like Claude, Cursor, or Windsurf without needing your own API keys.

A+ Quality Score 96.67/100

new-york

nyc

open-data

socrata

soda

catalog



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.

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.

NYC Open Data Explorer MCP

5 tools available

Cloud-hosted on Vinkius

You no longer need to manually download CSVs or navigate complex web portals to find city information. This MCP acts as a direct interface to the Socrata SODA API used by the City of New York. Your agent can now browse the full catalog of 3,000+ datasets, ranging from 311 service requests to traffic patterns and building permits.

Instead of just looking at a static list, your AI can actively investigate. It can find the right dataset using keywords, check exactly which columns are available, and then pull specific rows using SOQL filters. If you need to know the scale of a dataset before running a heavy query, your agent can check the total row counts or get grouped statistics for specific columns. It's a way to turn the city's massive public data repository into a live, searchable knowledge base for your AI-driven analysis.

Core Capabilities

01 — Catalog Discovery

Your agent searches through 3,000+ datasets to find relevant information using keywords.

02 — Metadata Inspection

The AI checks column names and descriptions to ensure queries are formatted correctly.

03 — Filtered Row Retrieval

Your agent pulls specific data points using SOQL syntax for precise filtering.

04 — Data Aggregation

The AI calculates row counts and grouped statistics to summarize large datasets.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/nyc-open-data-explorer — connect your AI agent in three steps.

- 01 Connect your MCP-compatible client to the Vinkius platform.
- 02 Ask your AI client to find a specific dataset using keywords.
- 03 Use the dataset ID to inspect the available columns.
- 04 Run a filtered query to pull the specific rows you need.

Connect your client to Vinkius and start querying NYC data immediately.

Built For

Data professionals and researchers who need to pull real-time city information into their AI workflows.

Data Analysts

They hand off the task of finding and filtering specific city datasets to their agent.

Urban Researchers

They use the MCP to pull statistical summaries of city trends like crime or traffic.

Software Developers

They integrate live NYC data into applications by letting their AI client query the SODA API.

What Changes When You Connect

- 01 Eliminates the need for manual API key management or local data downloads.
- 02 Provides direct access to the official Socrata SODA API used by NYC.
- 03 Allows for precise data filtering using standard SOQL syntax.

04 Enables rapid dataset discovery through keyword searching.

Real-World Applications

Real-time Trend Analysis

An analyst asks the AI to find the most recent 311 complaints in a specific borough.

Automated Reporting

An agent pulls grouped statistics on traffic patterns to populate a report.

Dataset Auditing

A researcher checks the last update time and row counts to verify data freshness.

Schema Discovery

A developer uses the MCP to find the exact column names needed for a custom query.

Patterns to Avoid

Blind querying without schema checks

❌ AVOID

You ask your agent to pull specific columns from a dataset without knowing if those columns actually exist.

✓ INSTEAD

Always use ``get_dataset_info`` first to pull the exact column names and metadata before attempting to query.

Massive data pulls without sizing

❌ AVOID

You tell your agent to grab every row from a massive 311 call dataset to perform an analysis.

✓ INSTEAD

Check the scale of the data first using ``dataset_row_count`` to see if you need to apply filters or limit your results.

Vague keyword searches

❌ AVOID

You try to find data by typing generic terms like 'city info' into the search tool.

✓ INSTEAD

Use specific keywords like 'potholes' or 'taxi trips' with ``search_datasets`` to get relevant dataset IDs.

The Right Fit

Connect this MCP if you need to pull real-time NYC public records into an agentic workflow without managing API keys. Do not connect it if you only need static, historical CSV files. The decision depends on whether you need your AI to query live city data on demand.

5 Tools for Navigating NYC Open Data

These tools give your agent direct access to the Socrata SODA API to find, inspect, and query over 3,000 city datasets.

#	TOOL	DESCRIPTION
01	browse_dataset	Query specific rows from any NYC Open Data dataset using SOQL filters. You can use double quotes for string values and ISO format for dates.
02	dataset_row_count	Get the exact total number of rows in a dataset. You can apply filters to get a count for a specific subset of data.
03	dataset_stats	Retrieve row counts and grouped statistics for a dataset. This is useful for seeing the top-N values within a specific column.
04	get_dataset_info	Fetch the metadata for a single dataset. Use this to find the exact column names and descriptions before running queries.
05	search_datasets	Search the entire NYC Open Data catalog using keywords. It returns the dataset ID, name, and description for each match.

See It in Action

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

U Which NYC datasets are about potholes?



`search_datasets("potholes")` returns the matching datasets like DOT Pothole Work Orders and 311 Street Condition requests with their IDs.

U What columns does the 311 service requests dataset have?



`get_dataset_info("erm2-nwe9")` provides the full column list including `complaint_type`, `descriptor`, `borough`, and `incident_address`.

U How big is the 311 dataset and what are the top complaint types?



`dataset_row_count("erm2-nwe9")` shows the total rows, while `dataset_stats` with a `group_by` on `complaint_type` shows the most frequent complaints.

Frequently Asked Questions

01 Do I need an API key to use this MCP?

No. All access is anonymous and requires no credentials or API keys.

02 How do I know which columns to query?

You should use the `get_dataset_info` tool first to see the exact column names and descriptions.

03 What happens if I want to filter by a specific borough?

You can use the `browse_dataset` tool with a SOQL filter, such as `borough="QUEENS"`.

04 Can I get a summary of the data without downloading it?

Yes, you can use `dataset_stats` to get grouped counts or `dataset_row_count` for total totals.

05 Which AI clients can use this?

Any MCP-compatible client like Claude, Cursor, or Windsurf can use this MCP once connected via Vinkius.

06 Do I need an API key?

No. NYC Open Data serves all public datasets anonymously over its SODA API (data.cityofnewyork.us). This MCP defines no credentials and needs nothing configured.

07 How do I find the id of a dataset?

Use `search_datasets` with a keyword — it queries the platform catalog and returns each match's 8-character id (like "erm2-nwe9"). That id feeds `get_dataset_info`, `browse_dataset`, `dataset_stats` and `dataset_row_count`.

08 Borough values differ across datasets — how do I know the format?

Formats are dataset-specific: 311 and shootings use UPPERCASE ("QUEENS"), food and sanitation use Title Case ("Queens"), DOT potholes use single-letter codes. Tool descriptions state the expected values, and the errors from a mismatched value are reported verbatim by the platform.

09 What size is the catalog?

The platform catalogs about 3,000 entries: roughly 2,400 standard datasets, plus filtered views, maps and external data sources. `search_datasets` searches all of them; the `ny-*` suite MCPs wrap the most-used families with dedicated tools.

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>"nyc-open-data-explorer": { "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

NYC Open Data Explorer 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 NYC Open Data Explorer. 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	September 2026
MCP Server	NYC Open Data Explorer MCP
Server ID	01a0dfce-bb90-7352-af8d-28231bdd983a
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/en/ai-agent-connect/nyc-open-data-explorer.