

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

Amsterdam Open Data Explorer MCP.

Give your agent direct access to Amsterdam's live data platform.

Amsterdam Open Data Explorer MCP acts as a direct interface to the City of Amsterdam's DSO Dataplatform. Your AI client can discover over 122 datasets, inspect API specifications, and query live data tables without needing an API key. It turns the city's entire data catalog into a searchable, actionable resource for your agent.

A+ Quality Score 98.33/100

amsterdam

open-data

metadata

explorer

catalog

dso-dataplatform



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.

Amsterdam Open Data Explorer MCP

5 tools available

Cloud-hosted on Vinkius

You can now let your agent navigate Amsterdam's open data ecosystem as if it were a local database. Instead of manually searching through web portals, you use this MCP to let your AI client crawl the DSO Dataplatform. It can find specific datasets by keyword, pull the exact OpenAPI specs for any resource, and fetch actual rows of data using standard filters and pagination. Whether you are looking for specific water management stats or checking the city's data product catalog, this MCP provides the bridge between your agent and the city's live information. It handles the heavy lifting of navigating the platform's structure, so your agent can focus on analyzing the information it finds.

Core Capabilities

01 — Dataset Discovery

Your agent finds relevant datasets using keywords or tags.

02 — API Inspection

The agent reads OpenAPI specs to understand how to interact with specific resources.

03 — Live Data Fetching

Your agent pulls actual records from city tables using filters and pagination.

04 — Catalog Browsing

The agent identifies data products and the teams responsible for them.

One Click on Vinkius — From Prompt to Execution

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

- 01 Connect your MCP-compatible client to Vinkius.
- 02 The agent uses `list_datasets` to find relevant information.
- 03 The agent calls `get_dataset_info` to understand the API structure.
- 04 The agent uses `browse_dataset` to pull the specific data rows needed.

Connecting to this MCP gives your agent immediate access to Amsterdam's data infrastructure.

Built For

This MCP is built for anyone building agents that need to interact with municipal data without manual searching.

Data Scientists

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

Software Developers

They use the agent to prototype applications against live city data.

Urban Researchers

They use the agent to query specific city metrics and data products.

What Changes When You Connect

- 01 No API keys are required to access the DSO platform.
- 02 The agent discovers dataset structures automatically via OpenAPI specs.
- 03 Direct table access allows for precise data filtering and sorting.

- 04 The entire city data catalog is accessible through a single connection.
-

Real-World Applications

Automated Data Discovery

An agent searches for all datasets related to 'transportation' to build a local knowledge base.

Real-time Data Analysis

A researcher queries the latest water level datasets to check for specific environmental trends.

API Prototyping

A developer uses the agent to test how different dataset resources respond to specific queries.

Organizational Mapping

A user identifies which city data teams own specific products for contact or inquiry purposes.

Patterns to Avoid

Manual Portal Searching

❌ AVOID

You spend an hour clicking through web portals to find a specific water management CSV file.

✓ INSTEAD

Let your agent find the exact dataset using keywords with `'list_datasets'` instead.

Guessing API Endpoints

❌ AVOID

You try to guess the URL structure for a dataset to pull records.

✓ INSTEAD

Use `'get_dataset_info'` to read the OpenAPI specs so your agent knows the exact resources available.

Overwhelming the Agent with Raw Data

❌ AVOID

You ask your agent to read every single row in a massive city table at once.

✓ INSTEAD

Use `'browse_dataset'` to apply filters and pagination so the agent only gets the specific rows it needs.

The Right Fit

Connect this MCP if you need to build agents that query municipal data without manual searching. Do not connect it if you only need static, non-city datasets. The decision depends on whether your workflow requires live access to Amsterdam's DSO Dataplatform.

5 Tools for Navigating Amsterdam Open Data

These tools let your agent crawl the DSO Dataplatform, inspect API specs, and pull live records from city tables.

#	TOOL	DESCRIPTION
01	list_datasets	This tool finds all open datasets published by Amsterdam. You can filter the results using specific keywords or tags.
02	get_dataset_info	This tool retrieves metadata and OpenAPI specifications for a single dataset. It shows you exactly which resources and operations are available.
03	list_data_products	This tool accesses the city's catalog of official data products. You can use keywords or pagination to find what you need.
04	browse_dataset	This tool fetches actual data from any dataset resource. You can apply filters, sort results, or use pagination and field projection.
05	list_data_teams	This tool identifies the specific city data teams that own the various data products.

See It in Action

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

U Which Amsterdam datasets are about water?



The `list_datasets` tool with a 'water' keyword returns matching services, including their descriptions, tags, and API URLs.

U What resources does the bomen dataset expose?



Using `get_dataset_info` for 'bomen' provides the OpenAPI spec and lists available resources like `stamgegevens` and `kapenherplant`.

U Fetch the first 10 rows of bag nummeraanduidingen.



The `browse_dataset` tool retrieves the first 10 address rows from the 'bag' dataset using the 'nummeraanduidingen' resource.

Frequently Asked Questions

01 Do I need an API key for Amsterdam open data?

No, this MCP allows you to explore and query the DSO platform without an API key.

02 How many datasets can I access?

You can access all 122+ datasets published on the Amsterdam DSO Dataplatform.

03 Can I filter the data I retrieve?

Yes, the `browse_dataset` tool supports filters, sorting, pagination, and field projection.

04 Which AI clients can use this MCP?

You can use this with any MCP-compatible client like Claude, Cursor, Windsurf, or VS Code.

05 How does the agent know which API endpoints to use?

The agent uses the `get_dataset_info` tool to read the OpenAPI specification for any given dataset.

06 Do I need an API key?

No. The DSO dataset index, the per-dataset OpenAPI specs and the generic table access are public (OPENBAAR). Some individual datasets are marked as requiring a key in their index entry — the dataset info response shows the authentication requirement, so check it before assuming access.

07 How do I find the right resource names?

Call `get_dataset_info` for the dataset: it fetches the live OpenAPI spec and lists the exposed resources and their operations. Use the resource name (usually the Dutch word) with `browse_dataset` — e.g. 'stamgegevens' for bomen, 'nummeraanduidingen' for bag.

08 What does `browse_dataset` return?

A page of the dataset's rows: the items array, current page number and page size, plus a next-link (next) when more pages exist. Filters are passed as comma-separated field=value pairs and only applied when the dataset actually supports that filter field.

09 What is the difference between the catalog and the datasets?

The dataset index (`list_datasets`) covers the machine-readable REST APIs. The data catalog (`list_data_products` / `list_data_teams`) covers the city's curated catalogue of data products and the 26 data teams that produce and own them — useful for mapping a question to the right dataset or owner.

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>"amsterdam-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

Amsterdam 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 Amsterdam 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	Amsterdam Open Data Explorer MCP
Server ID	01a0d9c8-7dca-70e3-a1d8-bf7ac00a3f6b
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/amsterdam-open-data-explorer.