MCP SERVER | NO CODE | CLOUD HOSTED

Nasdaq Data Link MCP

Access structured financial data via natural language queries.

Nasdaq Data Link (Quandl) connects your AI agent to professional-grade financial and economic datasets. Use natural language to query datatables, check dataset schemas for metadata, or trigger massive bulk downloads from top vendors like the S&P 500 index providers. It lets you pull complex financial history without writing a single API script.

A+ Quality Score 98.33/100

financial-data | stock-market | economic-indicators | quandl | nasdaq

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/nasdaq-data-link-quandl

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.

Nasdaq Data Link (Quandl) MCP

4 tools available
Cloud-hosted on Vinkius

This MCP gives your agent direct access to professional-grade data sets covering finance and economics. Instead of spending hours figuring out which vendor table holds the fundamental data you need, you just ask for it. Your agent handles the connection and pulls unsorted records using specific filters like tickers or date ranges. Need to know what columns are even available before running a query? You can check the metadata first. For massive analysis, you don't have to handle pagination manually; simply request bulk downloads for entire datasets in formats like CSV or Parquet. This capability means you get raw data into your environment fast, letting you focus on modeling instead of data plumbing. All 4,000+ MCPs are available through Vinkius, making this a single place to access global financial intelligence.

Core Capabilities

01 — Query specific datasets

Tell your agent which datatable and filters you need, and it pulls the unsorted data directly.

02 — Inspect dataset structure

Check a table's metadata to see exactly what columns are available and how they should be filtered before querying.

03 — Initiate large exports

Request the system start generating massive files for entire datasets, receiving status updates like PENDING or RUNNING.

04 — Download processed files

After a bulk export finishes, you can download the final data file in CSV, Parquet, or ZIP format.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/nasdaq-data-link-quandl** — connect your AI agent in three steps.

01 Subscribe to this MCP and enter your Nasdaq Data Link API Key into your preferred AI client.

02 Use natural language prompts (e.g., 'Get the metadata for XYZ datatable') to let your agent inspect available datasets and required filters.

03 Execute a query or request a bulk download, and your agent delivers the raw data points or monitors the export status until you can download the final file.

The bottom line is that it lets you access complex financial history using simple conversation instead of custom code.

Built For

Anyone who spends time gathering historical, structured data—especially those whose jobs involve cross-referencing economic indicators or market prices from different sources. If your process involves writing complex API calls just to get a few columns of data, this MCP is for you.

Financial Analyst

Pulling historical price points or specific fundamental metrics (like revenue) across multiple years without manually adjusting date ranges in different spreadsheets.

Data Scientist

Exploring dataset schemas and running targeted bulk exports for model training, letting the agent handle the initial data plumbing work.

Quantitative Researcher

Automating the retrieval of economic indicators or fundamental data points to feed into a larger quantitative model input pipeline.

What Changes When You Connect

01 Stop writing complex API scripts. You can ask your agent to get unsorted data using specific filters (like tickers or dates) with the `get_datatable` tool, getting the numbers you need in seconds.

02 Before running a query, check the structure first. Use `get_datatable_metadata` to see exactly what columns exist and which ones can be filtered, preventing frustrating errors down the line.

03 Handling massive data sets is easy. Initiate huge exports for entire datasets using `request_bulk_download`. The agent tracks the status until you're ready to pull it with `get_bulk_download_file`.

04 The MCP handles large result set pagination and exporting, so you don’t have to worry about manually chunking data or dealing with API limits. Just tell your agent what you need.

05 You get access to professional-grade economic indicators and alternative financial metrics—data that is usually locked behind complex vendor portals.

Real-World Applications

Building a Quarterly Performance Report

A Financial Analyst needs Q3 data for 10 key stocks. Instead of running 10 separate API calls, they ask the agent to query datatables with advanced filtering for the specific quarter and list of tickers, getting all unsorted results in one go.

Modeling a Market Indicator Shift

A Quantitative Researcher wants to model how bond yields correlate with energy prices. They first use `get_datatable_metadata` to understand the available columns, then use `request_bulk_download` on both data sets for a decade's worth of records.

Investigating an Outlier Stock Price

A Data Scientist finds a strange price movement. They ask the agent to fetch the last 20 records from that stock’s datatable, instantly seeing the raw data points and confirming if the pattern is real or an error.

Preparing for Deep Research

A user needs a massive dataset of fundamental company metrics. They tell the agent to start the bulk download for the entire ZACKS/CP datatable, setting it up and getting status updates while they work on other tasks.

Patterns to Avoid

Assuming data structure

✗ AVOID

Writing a query to pull `unitid` when the table actually calls it `identifier`. The API returns an error, and you waste 30 minutes debugging column names.

✓ INSTEAD

First, use `get_datatable_metadata` on the specific datatable. This lets your agent confirm all available columns before you ask for data using `get_datatable`. It saves time.

Trying to pull a huge dataset in one go

✗ AVOID

Asking for 20 years of daily price history, which hits an API limit and only returns the first 500 records.

✓ INSTEAD

For large historical pulls, start by using `request_bulk_download`. This tells the system to handle the massive load in the background. You track its status until you can download it with `get_bulk_download_file`.

Manually cross-referencing multiple vendors

✗ AVOID

Having to switch between five different vendor portals, downloading CSVs, and manually joining the data in Excel.

✓ INSTEAD

Use your agent to query specific datatables from various sources sequentially. Then, use `get_datatable` for each one, gathering all necessary raw inputs into a single workflow.

The Right Fit

Use this MCP if your data needs are highly structured: you need time series pricing, fundamental financial metrics, or standardized economic indicators from known vendors. If the source is locked behind a vendor API and requires specific filtering (by date, ticker, etc.), this tool works. Don't use it if you need to analyze unstructured text—like reading earnings call transcripts or summarizing news articles. For that, you'd need an MCP focused on document processing or web scraping. If your data is in a database schema you control internally, using a direct SQL-based connector might be better than relying on vendor endpoints.

The Old Way: Data Assembly Hell

Right now, gathering comprehensive market intelligence is painful. You start by checking the S&P for historical prices in one tab, then jump to the FRED site to get economic indicators, and finally check a third vendor for alternative data—all requiring separate logins and manual downloads. You end up with three different CSVs that need hours of cleaning and joining before your analysis can even begin.

With this MCP, you tell your agent exactly what you need, like 'Give me the closing price for MSFT and the unemployment rate for the last five years.' The agent handles connecting to multiple sources and pulls all the raw data points into one stream. You get clean inputs; you skip the copy-paste hell.

Nasdaq Data Link: Accessing Professional Financial Metrics

You don't have to manually run through vendor documentation just to find out which column is actually filterable. Instead, you ask the agent to inspect the data structure using `get_datatable_metadata`. It tells you the exact schema and filtering capabilities right away.

The process changes from a multi-hour data plumbing project into a simple conversation. You simply tell your agent what question you need answered, and it retrieves the professional data needed instantly.

Nasdaq Data Link (Quandl) with 4 Tools

These four tools let you query specific financial datatables, inspect their structure, or request massive bulk exports directly from your AI agent.

#	TOOL	DESCRIPTION
01	`request_bulk_download`	Starts an export process for an entire dataset and reports the status (PENDING, RUNNING, SUCCEEDED) along with file URLs.
02	`get_bulk_download_file`	Downloads a specific bulk file once the data export has been successfully processed.
03	`get_datatable`	Pulls unsorted, filtered data points from a specified Nasdaq datatable using defined criteria like tickers and dates.
04	`get_datatable_metadata`	Retrieves the full description of a datatable, listing its column types and identifying which columns can be used as filters.

See It in Action

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

U Show me the metadata for the MER/F1 datatable.

I've retrieved the metadata for MER/F1. This table contains fundamental financial data. Available columns include 'reportid', 'unitid', 'indicatorid', and 'value'. The columns 'ticker' and 'date' are primary filters.

U Fetch the latest 5 records from WIKI/PRICES for ticker MSFT.

Querying WIKI/PRICES for MSFT... I found the most recent data points. On 2018-03-27, the close price was 89.47. Would you like me to export the full history for this ticker?

U Request a bulk download for the ZACKS/CP datatable.

I have initiated the bulk download request for ZACKS/CP. The current status is 'PENDING'. I will monitor the export and let you know when the file is ready for download.

Frequently Asked Questions

01 **How do I check which columns are available in a Nasdaq datatable?**

Use the `get_datatable_metadata` tool. This tells you the table's description, its column types, and confirms exactly which fields can be used for filtering before you run a full query.

02 **Can I download huge amounts of financial data with Nasdaq Data Link (Quandl)?**

Yes. If the dataset is too large to fetch directly, use `request_bulk_download` first. The system handles the heavy lifting in the background and lets you pull the final file using `get_bulk_download_file`.

03 **What if I only need a few rows of data?**

Use the `get_datatable` tool. Just provide the datatable name, your desired filters (like date range and ticker), and it will pull the unsorted results directly to your agent.

04 **Does this MCP work with all financial data?**

It accesses datasets from various vendors available through Nasdaq Data Link. You must specify the exact datatable name when making a query or metadata request.

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 `"nasdaq-data-link-quandl": { "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

Nasdaq Data Link (Quandl) 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 Nasdaq Data Link (Quandl). 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	Nasdaq Data Link (Quandl) MCP
Server ID	`019e38de-5af5-714f-9421-4539362341b4`
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/nasdaq-data-link-quandl.