MCP SERVER | NO CODE | CLOUD HOSTED

CData Connect Cloud MCP for AI Agents

Manage cross-system data sources and API schema mapping

CData Connect Cloud is a universal data gateway that lets your AI client read and execute complex queries across dozens of disparate systems. It dynamically maps structures, proxies native APIs, and parses SQL schemas without you having to write custom integration code for every source.

A+ Quality Score 100/100

data-integration | sql-proxy | api-gateway | schema-mapping | data-connectivity | unified-data

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/cdata-connect-cloud

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.

CData Connect Cloud MCP

8 tools available
Cloud-hosted on Vinkius

Dealing with modern data means dealing with a dozen different endpoints—a database here, an API there, maybe some legacy system somewhere else. Usually, getting clean data requires writing a whole new connection script just for that one source. That’s slow and painful.

CData Connect Cloud changes the game. It lets your AI client treat all those separate APIs like they're one connected database. You run a standard SQL query, and this MCP handles the complex routing and schema mapping behind the scenes. Your agent figures out how to talk to five different systems using one set of instructions.

If you’ve ever needed to pull data from multiple sources into one place without building middleware, this is it. When your AI client connects through Vinkius, they gain access to CData Connect Cloud alongside thousands of other specialized tools in the catalog, making complex cross-system queries manageable and reliable.

Core Capabilities

01 — Map Data Sources

List all external endpoints and databases that are already connected through your account.

02 — Test Connectivity

Check the connection health of any active data source to ensure the link is live and working before running a query.

03 — Inspect Schemas

Examine the full structure of an entire database or specific tables to know exactly what fields are available for querying.

04 — Execute Queries

Run a direct SQL query against any connected data source, pulling back structured results immediately.

05 — Build Connections

Programmatically establish and configure entirely new backend data proxies using the CData platform's logic.

One Click on Vinkius — From Prompt to Execution

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

01 First, you use the tool to check available connections or build a brand new connection proxy by providing credentials.

02 Second, your agent uses schema inspection tools to map out exactly which tables and columns are available across all connected endpoints.

03 Finally, your agent executes the query using the data source's unique identifier and the defined SQL logic, returning clean records.

The bottom line is that this MCP lets you write one set of instructions for complex queries spanning multiple systems without touching underlying connection code.

Built For

This MCP targets developers and architects who spend their time connecting data —the Data Engineers, the API Architects, and Integration Leads. If your job involves combining information from Salesforce with Postgres and some weird SaaS tool, you need this.

Data Engineer

You use this to pull unified datasets by querying multiple operational databases in a single, structured run.

API Architect

You leverage the connection tools to validate and map complex data flow paths between services before deployment.

Integration Platform Lead

You use it to build reliable, multi-source connectivity proofs of concept quickly, validating structural integrity across systems.

What Changes When You Connect

01 You don't have to write custom code for every source. This MCP lets your agent treat multiple APIs like one unified database, saving massive development time.

02 Validate connectivity instantly. Use cdata_test_connection to check a proxy's latency before running expensive queries, preventing failed jobs.

03 See everything available at once. Running cdata_list_connections or listing workspaces lets your agent map the entire data landscape for you.

04 Deeply understand data structures. Tools like cdata_get_schema_metadata let you programmatically explore every field and relationship in a system before querying it.

05 Keep things clean. Instead of pulling raw, unstructured blobs, running an execute query gives you structured records that your agent can use directly.

Real-World Applications

Combining CRM and Billing Data

A user needs to build a report combining customer names from the core CRM database with payment history from a separate billing API. Instead of building two separate data pipelines, the agent runs one query, leveraging cdata_execute_query across both sources.

Auditing Data Source Health

An architect needs to prove that all 15 connected systems are operational before a major migration. They use cdata_list_connections and then run multiple cdata_test_connection calls to validate the structural matrix.

Identifying Data Gaps

A data lead suspects some data is missing because they don't know what sources exist. They use cdata_list_workspaces and cdata_get_schema_metadata to audit all available logical scopes.

Automating Initial Setup

A new team member needs access to a legacy API that requires specialized credentials. Using cdata_create_connection, the agent programmatically establishes the required secure proxy link immediately.

Patterns to Avoid

Hardcoding connection strings

✗ AVOID

Writing explicit database names or API keys directly into the prompt for your AI client. This makes the system brittle and impossible to update when an endpoint changes.

✓ INSTEAD

Let your agent handle it by first listing available sources using cdata_list_connections, then defining scope via cdata_get_schema_metadata before finally executing the required data pull with cdata_execute_query.

Assuming schema visibility

✗ AVOID

Asking your agent to select a column name (e.g., 'CustomerStatus') when you haven't confirmed that column actually exists in the target database, leading to query failure.

✓ INSTEAD

Always check the available fields first by using cdata_get_table_columns on the relevant table before building your final SQL query.

Ignoring connection health

✗ AVOID

Running a complex data pull against a proxy that has been offline or is timing out, resulting in ambiguous failure messages and lost work.

✓ INSTEAD

Always start by running cdata_test_connection. This validates the structural matrix proxy latency before you commit to a full query.

The Right Fit

Use this MCP if your primary problem is combining data from disparate, non-standardized APIs or databases using SQL logic. You need programmatic access and schema mapping across multiple sources. Don't use it if you just need simple single-source querying, as a standard database connector will suffice. If the issue is managing credentials or user permissions across different systems, focus on cdata_list_workspaces for scoping data groups. Use this MCP when your workflow requires connecting and testing multiple independent endpoints; otherwise, stick to dedicated ETL tools.

CData Connect Cloud: Solving Multi-API Data Integration Pain

Manually pulling data from a company today means clicking through five different dashboards. You copy customer IDs from Salesforce, then open the billing system to get payment status, and finally jump into the support platform to find interaction history. Each step is manual, requires context switching, and you risk errors every time you copy-paste.

With CData Connect Cloud, your agent handles all those hops automatically. You give it one prompt: 'Get me a report on high-value customers who haven't paid in 90 days.' The MCP executes the necessary queries across the CRM, billing API, and support system, delivering the unified answer without you touching another UI.

CData Connect Cloud: Advanced Schema Mapping for Data Architects

Architects used to spend days writing boilerplate code just to discover what fields were available in a new system. They'd run exploratory scripts, manually checking schema definitions against documentation, slowing down the entire integration project.

Now, your agent can instantly map the environment. By running cdata_get_schema_metadata and cdata_list_tables, you get an immediate, comprehensive view of every logical data structure. It's not just faster; it changes how quickly you can prototype a new data layer.

CData Connect Cloud: 8 Tools for Data Integration and Schema Mapping

Use these tools to discover, test, configure, and query every external database source available through CData Connect Cloud.

#	TOOL	DESCRIPTION
01	`cdata_create_connection`	Sets up a completely new data source proxy using the CData system logic.
02	`cdata_execute_query`	Runs and routes an SQL query into any downstream database, returning clean values.
03	`cdata_get_schema_metadata`	Retrieves the full structural graph, showing every available data interaction mapped in the system.
04	`cdata_get_table_columns`	Shows detailed field definitions for a specific table boundary within a connection.
05	`cdata_list_connections`	Outputs an array listing every external data source that is currently connected through CData.
06	`cdata_list_tables`	Unpacks a list of available structural collections mapped securely within a given backend connection.
07	`cdata_list_workspaces`	Enumerates all logical virtual workspaces that segment different organizational data groups.
08	`cdata_test_connection`	Runs a ping check against a connected proxy to verify its operational status and latency.

See It in Action

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

U What are the latest connection status and what tables exist in my system?

System Connection Report

- ✅ **Production DB:** Operational (Latency: 45ms)
- ⚠️ **Billing API:** Warning - Check Credentials
- ✅ **CRM Data Store:** Operational (Latency: 32ms)

Available Tables:

- `workspaces.global_scope`: Contains organizational groupings.
- `connections.billing_data`: Includes tables like `payments`, `invoices`, and `accounts`.
- `tables.crm_users`: Holds user records, including `user_id` and `last_login`.

U Run a query for all users who logged in last month but haven't created an invoice.

Query Results: Unbilled Users

USER ID	NAME	LAST LOGIN DATE	STATUS
10045	Alice Johnson	2026-05-15	Needs Invoice
10987	Robert Lee	2026-05-01	Needs Invoice

Query executed successfully against the combined CRM and Billing data proxies.

U Can you list all available logical workspaces for me to check?

Workspace Catalog

The following distinct organizational scopes are currently mapped:

1. `workspaces.global_admin` : Master scope for system-wide settings.
2. `workspaces.marketing_campaigns` : Data specific to Q3 ad campaigns.
3. `workspaces.internal_hr` : Limited access data for HR records.

Frequently Asked Questions

01 **How does CData Connect Cloud help me query multiple databases at once?**

It treats separate APIs and databases as one virtual system. You write a single SQL instruction, and the MCP handles routing that command to all necessary sources, pulling together a unified data set.

02 **I need to check if my new API connection is actually working before I use it.**

You can run a quick test ping using the connectivity tools. This validates the link's health and measures latency, so you know your data source is active and reliable.

03 **What if my company adds a new database system I need to query?**

You can programmatically build it using connection tools. The MCP establishes the necessary secure proxy link, making the new system available for querying within minutes.

04 **Can CData Connect Cloud help me understand what data fields are in a table?**

Yes. By inspecting schemas and listing columns, you can see every field defined on any connected source before writing a single line of query code. This prevents errors.

05 **Is CData Connect Cloud only for large enterprise data systems?**

No. It's built to handle anything—from small, departmental databases to massive corporate API backends. The gateway adapts to whatever you connect it to.

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 `"cdata-connect-cloud": { "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

CData Connect Cloud 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 CData Connect Cloud. 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	CData Connect Cloud MCP
Server ID	`019d756a-cc4e-707b-a415-0c45b575390f`
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/cdata-connect-cloud.