

MCP SERVER

NO CODE

CLOUD HOSTED

NCDC Climate Data Online MCP

Query Decades of Historical Weather Records

NCDC Climate Data Online connects your AI agent directly to the National Climatic Data Center's massive archive of historical weather records. You can pull precise temperature and precipitation data, find specific global weather stations, or map out entire regions using natural conversation. Stop navigating complex government APIs; just ask for the climate history you need.

A+ Quality Score 100/100

climate-data

historical-weather

meteorology

environmental-monitoring

data-archiving



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.

NCDC Climate Data Online MCP

10 tools available

Cloud-hosted on Vinkius

Need to analyze how a certain location handled extreme heat back in 1985? This MCP lets your agent do that by talking directly to decades of authoritative NOAA data. Instead of manually searching through multiple databases and figuring out which station ID means what, you just tell the system what you need—a specific dataset for a date range or all stations within a county. The connector handles the complex querying, letting you focus on the science. When you connect this MCP via Vinkius, your agent gets instant access to over 4,000 other tools and can start building real-world climate models immediately.

Core Capabilities

01 — Retrieve historical weather records

Get actual daily or monthly temperature and precipitation measurements for a specific location and date range.

03 — Locate weather stations by geography

Search for specific physical monitoring sites worldwide, retrieving their full metadata and coverage details.

05 — Determine available measurement types

List specific scientific variables, such as Max Temperature or Snowfall, that the system tracks for a region.

02 — Find and inspect data sets

List available climate datasets, like GHCND or GSOD, to understand what historical records exist globally.

04 — Filter data by location type

Browse structured location categories (like State or Country) to narrow down your geographical search boundaries.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/ncdc-climate-data-online-1 — connect your AI agent in three steps.

- 01** First, subscribe to this MCP and provide your NCDC (NOAA) API Token. You'll need to grab this token from ncdc.noaa.gov.
- 02** Next, connect the MCP via your preferred client like Cursor or Claude. Your agent now has direct access to all historical data tools.
- 03** Finally, ask your AI client a natural language question—for instance, 'What was the average temperature in Asheville for January 2023?'—and it executes the necessary queries.

The bottom line is that you get direct access to decades of official climate records without writing complex API calls or dealing with multi-step data pipelines.

Built For

Anyone who works with environmental history, large datasets, or risk assessment. This hits the researchers needing deep historical proof and the analysts running models on real-world climate trends.

Climate Scientist

Uses this MCP to retrieve detailed records for climate modeling, comparing current data against decades of past weather patterns.

Environmental Data Analyst

Automates the gathering of historical environmental trends—like rainfall or max temperature—for reports on conservation or resource management.

Risk Assessor (Insurance/Agriculture)

Checks regional climate records to quantify historical risk, such as drought frequency or flood severity, for policy underwriting or crop planning.

What Changes When You Connect

-
- 01 Instantly model trends: Use `get_climate_data` to pull precise, historical records for temperature and precipitation without writing complex SQL or API calls.

 - 02 Understand the data source: If you don't know which dataset to use, run `list_datasets` first. It gives you a clear map of GHCND vs. GSOD.

 - 03 Target specific points: Need info on one station? Use `get_station` to pull all metadata, ensuring your query uses the right identifier.

 - 04 Drill down variables: Running through `list_data_categories`, followed by `list_data_types` lets you narrow a general query (like 'weather') down to exactly what you need (like 'Snowfall').

 - 05 Filter locations easily: Use `list_location_categories` and `list_locations` together. This prevents vague queries and locks your data to the correct city, state, or country.

 - 06 Start simple: You don't have to learn 10 tools at once. Start by asking your agent for a list of stations (`list_stations`) and see where it takes you.
-

Real-World Applications

Evaluating historical drought risk in the Midwest.

An agronomist asks their agent: 'What was the average precipitation for Iowa between 2010 and 2015?' The agent uses `get_climate_data` to retrieve the records, allowing the agronomist to calculate regional yield volatility.

Determining climate variability for insurance claims.

An analyst needs to compare Hurricane Season 2018 data against 1958. They use `list_datasets` to confirm the right dataset and then execute multiple calls to `get_climate_data`, generating a comparative report.

Writing a thesis on urban heat island effect.

A student asks: 'Find all weather stations in Chicago's downtown area.' The agent uses `list_stations` and then the location filters to scope the search, providing metadata needed for the academic paper.

Validating data coverage for a new monitoring project.

A researcher asks: 'What are all available data classes in this region?' The agent uses `list_data_classes` and then checks specific variables using `list_data_types` to ensure the system tracks required metrics.

Patterns to Avoid

Assuming a tool exists**X AVOID**

The user just asks, 'Give me all weather data for my area.' The agent fails because it doesn't know what specific variables or location IDs to use.

✓ INSTEAD

Don't ask for everything at once. First, run `list_location_categories` to narrow your geography, then check `list_data_types` to confirm the variable (e.g., 'Max Temperature'), and finally call `get_climate_data`.

Confusing dataset types**X AVOID**

The user requests data for a specific date range but doesn't know if they need a Daily or Monthly summary, leading to an incorrect query failure.

✓ INSTEAD

Before asking for records, run `list_data_classes` to understand the temporal flexibility of the data. This ensures you select 'Daily' when you want daily metrics.

Missing location context**X AVOID**

The user only inputs a generic state name into the query, and the agent returns ambiguous or irrelevant results.

✓ INSTEAD

Always run `list_stations` or use `list_locations` first. This grounds your request in actual station IDs, guaranteeing accurate records via `get_climate_data`.

The Right Fit

Use this MCP if your work requires historical, authoritative weather data sourced from NOAA's NCDc archive. You need to compare climate metrics across years or different geographical points, and you are dealing with established scientific variables like GHCND or GSOD. Don't use it if you just need a simple, single-point forecast for tomorrow; that requires a live forecasting tool. Also, don't use it

if your data comes from an internal company database; in that case, a dedicated internal connection will work better. This MCP is purely for external, historical, public climate record retrieval.

Gathering Historical Climate Data Used to Be a Nightmare

If you've ever needed weather data from ten years ago, you know the drill. You have to navigate NOAA websites, figure out which specific dataset ID corresponds to what you want, and then manually piece together station IDs for every city on your list. It involves downloading multiple CSVs, renaming columns, and cross-referencing dates just to get a basic trend line.

With this MCP, the process collapses into a single conversation with your agent. You describe the problem—'I need temperature data for three counties in 2015.' Your agent handles listing the necessary stations (`list_stations`), confirming the correct variables (`list_data_types`), and executing `get_climate_data` across all of them, giving you a clean output ready for analysis.

Get Climate Data with NCDC Climate Data Online

You no longer need to jump between five different government portals just to find one data point. The MCP automatically handles the complexity of location intelligence, ensuring that when you ask for a county, it uses the correct internal identifiers.

What's different now is speed and reliability. You get immediate access to structured, authoritative climate history—all through your AI client.

NCDC Climate Data Online: 10 Tools

These tools allow your agent to list, find, and retrieve specific climate datasets, station metadata, and historical weather records from the NCDC archive.

#	TOOL	DESCRIPTION
01	<code>get_climate_data</code>	Pulls actual climate data records based on specified location, date range, and metric.
02	<code>get_dataset</code>	Retrieves detailed information about a specific climate dataset by its identifier.
03	<code>get_station</code>	Fetches comprehensive metadata, coverage area, and details for one particular weather station.
04	<code>list_data_categories</code>	Lists the high-level scientific categories of data available in the system (e.g., Temperature, Precipitation).
05	<code>list_data_classes</code>	Shows the frequency classes for available data, such as Hourly, Daily, or Monthly summaries.
06	<code>list_data_types</code>	Lists specific measurable variables within a category, like Max Temperature or Snowfall.
07	<code>list_datasets</code>	Provides a list of major NCDC climate datasets available for querying (e.g., GHCND).
08	<code>list_location_categories</code>	Shows the hierarchy of geographical locations, like Country or State.
09	<code>list_locations</code>	Lists specific physical locations that can be used to filter data queries.
10	<code>list_stations</code>	Provides a list of all weather monitoring stations and their unique IDs across the globe.

See It in Action

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

U List all weather stations in the city of Asheville, NC.



I've retrieved the weather stations for Asheville, NC. There are several active stations, including 'ASHEVILLE REGIONAL AIRPORT, NC US' (GHCND:USW00003812) and 'ASHEVILLE 13 S, NC US'. Would you like the detailed metadata for the airport station?

U Get daily temperature data for station GHCND:USW00003812 for January 2023.



Fetching climate records... For January 2023 at Asheville Regional Airport, the average maximum temperature was 52°F (11°C) and the minimum was 34°F (1°C). There were 15 days with recorded precipitation. Shall I provide the full daily breakdown?

U What climate datasets are available for global daily summaries?



Retrieving datasets... For daily summaries, the primary datasets are GHCND (Global Historical Climatology Network Daily) and GSOD (Global Summary of the Day). GHCND is generally recommended for its extensive station coverage. Would you like more info on GHCND?

Frequently Asked Questions

01 How do I find out what kind of weather data NCDC Climate Data Online can handle?

You use the `list_data_types` tool. This function gives you a list of all specific variables tracked by the system, such as Max Temperature or Snowfall, so you know exactly what to ask for.

02 What is the difference between GHCND and GSOD in NCDC Climate Data Online?

You can find out using `list_datasets`. These datasets are distinct archives; GHCND generally provides more detailed station coverage, while GSOD offers a different global summary format.

03 Can I get climate data for multiple stations at once with NCDC Climate Data Online?

Yes. First, use `list_stations` to find the IDs of all needed sites. Then, pass those IDs and your date range into a single query using `get_climate_data`.

04 What is required to use NCDC Climate Data Online MCP?

You must subscribe to the MCP and provide your personal NCDC (NOAA) API Token. This token grants your agent authenticated access to the official records.

05 How do I check what location types are available for data filtering?







Run the `list_location_categories` tool. This will show you if the system supports filtering by Country, State, County, or City, guiding your subsequent queries.

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>"ncdc-climate-data-online-1": { "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

NCDC Climate Data Online 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](https://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 NCDC Climate Data Online. 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	NCDC Climate Data Online MCP
Server ID	019d75db-9c09-7247-8e2e-693f0f47ff3a
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/ncdc-climate-data-online-1.