

MCP SERVER

NO CODE

CLOUD HOSTED

Stanford GDELT MCP

Map Global Media Events & Themes.

Stanford GDELT MCP gives you access to the world's largest open dataset for monitoring global news media coverage in real time. Track how stories spread across countries, languages, and political themes, analyzing everything from sentiment shifts to geographic hotspots. Use this MCP to turn raw global data into actionable intelligence for journalists, policy makers, and PR teams.

A+ Quality Score 100/100

gdelt

global-news

sentiment-analysis

media-monitoring

geopolitics

tv-news



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.

Stanford GDELT MCP

16 tools available

Cloud-hosted on Vinkius

This connector links your AI client directly to the GDELT Project API, giving you a direct line to billions of records of global news coverage. You can monitor what's happening worldwide by analyzing articles across 100+ languages, tracking specific themes like climate change or elections, and seeing exactly where media attention is focused geographically. It's far more than just searching; it helps you map the story itself—who reported it, how they felt about it (the tone), and which nations are paying attention. If you're relying on manual dashboards or basic search engines to track global events, this MCP changes that. Through Vinkius, you connect your agent once and gain instant access to this massive data catalog, letting your AI client do the heavy lifting of complex cross-cultural analysis for you.

Core Capabilities

01 — Map Global Coverage

Pinpoint news events and measure media focus using coordinates or heatmap overlays.

03 — Identify Thematic Hotspots

Classify news content using hundreds of standardized themes covering politics, health, and economics.

05 — Monitor TV vs. Online Narrative Shifts

Compare how stories are covered in closed-captioned television news versus general online articles.

02 — Analyze Topic Trends Over Time

Track how the volume, sentiment, or language coverage of a specific subject changes over months or years.

04 — Deep Dive into Language/Region Spread

Determine which countries or linguistic communities are generating the most coverage on a topic.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/stanford-gdelt — connect your AI agent in three steps.

- 01** Subscribe to this MCP on Vinkius, which connects your AI client to the entire GDELT Project API.
- 02** Your agent accesses the structured tools and functions necessary for global data queries. No API key is needed because the connection is managed by Vinkius.
- 03** You prompt your AI client with a complex request—like 'Show me how sentiment around AI regulation changed in Europe over 6 months'—and the MCP executes the required steps across multiple data points.

The bottom line is, you tell your agent what global insight you need, and it handles the complex data calls to retrieve the precise information.

Built For

Journalists and political analysts who get tired of sifting through siloed dashboards.

If you need to know **why** a story is breaking now—and where its attention originated—you need this.

Investigative Journalist

Needs to track the initial reporting of an event, tracing coverage from small local papers to major international outlets by using `search_articles` and `get_geo_data`.

Crisis Communications Specialist

Must monitor brand reputation in real time. They use `get_timeline_tone` and `get_tone_chart` to track sentiment shifts across different regions after a major announcement.

Policy Analyst

Analyzes geopolitical discourse, using `search_by_theme` and `get_themes` to understand how standardized topic categories (like `TERROR` or `PROTEST`) are being covered globally.

What Changes When You Connect

-
- 01** Track the life cycle of a story. Instead of just seeing one article, you use `get_timeline_volume` to see exactly when media attention spikes and how long that coverage lasts.

 - 02** Understand global context instantly. Use `search_by_country` or `search_by_language` to filter results immediately, letting you compare US vs. UK coverage on the same issue.

 - 03** Measure emotional impact precisely. `get_timeline_tone` tracks sentiment shifts over time, so you know if public opinion is cooling down or heating up around a topic.

 - 04** Go beyond keywords. Use `search_nearby` when you need to find articles where two specific terms appear close together, which is much more accurate than a basic keyword query.

 - 05** Compare media types. You can run `get_tv_timeline` and compare the volume of coverage on broadcast news versus what's available in general online archives.
-

Real-World Applications

Tracking Political Fallout from an Election

A political analyst needs to know how different regions viewed a policy change. They use `search_by_theme` for 'ELECTION' and then run `get_timeline_country` to visualize which nations escalated their coverage and when that interest became global.

Researching Environmental Policy Spread

A policy expert wants to see how climate concerns are framed. They use `search_by_theme` for 'ENV_CLIMATECHANGE' and then utilize `get_timeline_lang` to compare coverage differences between English, Mandarin, and Spanish sources.

Monitoring Brand Reputation During a Crisis

A PR manager needs immediate insight. They use `get_tone_chart` on the company name over the last 3 months, seeing if negative sentiment spikes correlate with specific geographic areas identified by `get_geo_data`.

Investigating Complex Scientific Claims

A journalist needs to find nuanced reporting. Instead of just searching for 'vaccine side effects,' they use `search_nearby` with the terms 'vaccine' and 'adverse event' plus a distance parameter, ensuring the context is right.

Patterns to Avoid

Treating it like Google Search

✗ AVOID

A user searches for 'China economy' in their AI client and expects a list of recent headlines. They get raw data, but no context on how the story is spreading.

✓ INSTEAD

Use `search_articles` with filters to narrow down the scope (e.g., only from China) *and* follow up by running `get_timeline_country` or `get_themes` to give that headline historical and contextual depth.

Focusing only on keywords

✗ AVOID

A user searches for 'conflict' and gets thousands of articles, but can't tell if the tone was positive (de-escalating) or negative (violence).

✓ INSTEAD

Always follow a general search with `get_timeline_tone`. This gives you the emotional layer—the actual feeling behind the words—which is crucial for analysis.

Missing geographic context

✗ AVOID

A user sees an article about 'water scarcity' and assumes it applies everywhere, missing which specific nations are most affected or reporting on it.

✓ INSTEAD

Always run `get_geo_data` immediately after a search. This provides coordinates and location names, helping you pinpoint the exact source of the media attention.

The Right Fit

Use this MCP if your core problem is measuring global scale or tracking how *attention* moves—whether that's political focus, emotional tone, or journalistic coverage. You need to know not just *what* was said, but *where*, *when*, and *by whom*. Don't use this if you simply need basic keyword searching; those simple searches are better handled by generic web search tools. Similarly, don't rely on it if your goal is only localized data (e.g., just California news). If you're restricted to one country or one language for the entire analysis, then a simpler, local API might suffice. But because you need cross-border comparison—comparing US coverage of 'migration' with EU coverage of 'migration'—this MCP's breadth is what you require.

Mapping Global Attention Spreads Is Hard

Right now, tracking a single story across the world means jumping through endless dashboards. You open one tab for US news, another for Chinese media reports, and yet another to check sentiment via a separate dashboard. Copying dates, filtering languages, and trying to manually correlate how coverage shifts from online articles to broadcast TV is tedious at best.

With this MCP, your agent handles the entire process. You give it the topic, and it executes multiple data calls—running `get_timeline_volume` for spikes, then `getting_geo_data` for locations, and finally using `search_by_theme` to categorize the whole narrative. You get a single, synthesized answer, not three dozen links.

Analyze Coverage With The GDELT Theme Distribution

Manually analyzing themes across different articles is nearly impossible. You'd have to read hundreds of headlines and manually categorize them into 'ECONOMICS,' 'HEALTH,' or 'PROTEST.' This requires specialized human labor just for categorization.

Now, you use `get_themes`. Your agent runs the classification against massive datasets, instantly telling you that 40% of the coverage falls under 'ECONOMIC_BANKRUPTCY' and 25% is tagged as 'HEALTH_PANDEMIC.' You skip the reading and go straight to the insights.

Stanford GDELT: 16 Tools for Global Data

These tools let your agent perform highly specific data analysis, from mapping geographical hotspots to tracking sentiment changes over time.

#	TOOL	DESCRIPTION
01	<code>get_geo_data</code>	Retrieves geographic point data for news events, showing coordinates and location names for mapping.
02	<code>get_tv_channels</code>	Lists available TV news channels to understand the scope of broadcast analysis.
03	<code>get_tv_timeline</code>	Tracks how much airtime specific stories receive on television compared to general online articles.
04	<code>get_themes</code>	Retrieves standardized GDELT themes, classifying news content across broad topics like politics or health.
05	<code>get_timeline_country</code>	Shows the source country distribution timeline to identify which nations are covering an issue and when global interest grows.
06	<code>get_timeline_lang</code>	Provides a language distribution timeline, showing how different linguistic communities cover a topic over time.
07	<code>get_timeline_tone</code>	Generates a sentiment and tone timeline to track public opinion shifts related to any given subject.
08	<code>get_timeline_volume</code>	Provides the news volume timeline, which tracks overall media attention spikes and helps understand a story's lifecycle.
09	<code>get_tone_chart</code>	Generates a tone distribution chart to visualize if coverage is mostly positive, negative, or neutral.
10	<code>get_word_cloud</code>	Extracts the dominant key terms and concepts associated with a topic in media coverage.
11	<code>search_articles</code>	Searches global news articles using filters for date range, language, country, or general keywords.
12	<code>search_by_country</code>	Narrows the search results to only include news published from a specific country code.
13	<code>search_by_language</code>	Limits searches to articles written in one of over 100 specified language codes.

#	TOOL	DESCRIPTION
14	search_by_theme	Filters news results based on standardized topic categories like ENVIRONMENT or ELECTION.
15	search_nearby	Searches for articles where two specific terms appear near each other, providing a more precise search than keyword matching.
16	search_tv	Finds and analyzes transcripts from major television news networks using keywords.

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 news articles about AI regulation?



I've searched GDELT for the latest articles on "AI regulation" from the past week. Coverage spans multiple countries with articles from the US, EU, UK, and China discussing frameworks, legislation, and safety standards.

U How has sentiment about climate change evolved over the last 3 months?



I've retrieved the tone timeline for "climate change" over the past 3 months from GDELT. The sentiment shows fluctuations correlating with major climate events and policy announcements.

U Search for TV news clips mentioning quantum computing



I've found TV news clips mentioning "quantum computing" from GDELT's TV news archive. Clips include segments from CNN, MSNBC, and Fox Business discussing breakthroughs and investment in quantum technology.

Frequently Asked Questions

01 Can I track sentiment over time using Stanford GDELT MCP?

Yes. Use `get_timeline_tone`. This tool tracks how public opinion changes around a topic, providing positive and negative values so you can monitor shifts in brand reputation or political favor.

02 Does Stanford GDELT MCP support keyword searching across multiple languages?

Yes. You use `search_articles` by combining language codes (like 'english' and 'spanish') with a date range to gather global coverage on the same issue.

03 How do I find out which country is most interested in an issue?

Use `get_timeline_country`. This tool analyzes source country distribution, showing you exactly when and where media attention for a given topic originates globally.

04 Is Stanford GDELT MCP only for major news outlets?

No. It covers global news sources across 100+ languages and includes specialized searches like `search_tv`, which pulls transcripts from networks such as CNN and BBC.

05 What is the best way to compare two topics in GDELT?







Start by using `get_timeline_volume` for both topics over the same time frame. This lets you graph their attention spikes against each other, revealing correlation or divergence.

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>"stanford-gdelT": { "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

Stanford GDELT 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 Stanford GDELT. 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	Stanford GDELT MCP
Server ID	019dea60-babd-7187-9a34-7f464919a360
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/stanford-gdelt.