

MCP SERVER

NO CODE

CLOUD HOSTED

Scopus MCP

Run advanced research and citation analysis.

Scopus MCP connects your AI client to Elsevier's massive academic literature and citation database. You can search for papers using advanced Boolean logic, build author profiles by checking H-indices, and calculate journal impact metrics like CiteScore—all from a natural conversation.

A+ Quality Score 100/100

academic-literature

citation-database

research-metrics

bibliometrics

metadata-search



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.

Scopus MCP

10 tools available

Cloud-hosted on Vinkius

You don't have time to manually log into multiple databases just to compile a literature review or audit an academic's output. This MCP lets you run advanced queries against the world's largest research index directly through your AI client. Instead of copy-pasting DOIs and running separate reports, you talk to your agent and it handles the complexity. You can find specific papers using Boolean syntax, pull full metadata for authors or institutions, or get citation overviews broken down by year.

Need to know how influential a journal is? Check its CiteScore or SJR metrics in one go. Want to see if a piece of research has social media buzz? The MCP pulls altmetrics data right into your workflow. When you connect this MCP via Vinkius, you get immediate access to these powerful tools without needing complex API calls or developer setup. It just works with natural language, letting you focus on the analysis, not the plumbing.

Core Capabilities

01 — Conduct deep literature searches

Search and retrieve metadata for academic abstracts using advanced search syntax.

02 — Build author and institution profiles

Fetch detailed professional profiles, including H-index counts and affiliation history for academics or universities.

03 — Analyze citation impact

Calculate comprehensive citation summaries over time, year by year, using a document's DOI or Scopus ID.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/scopus — connect your AI agent in three steps.

- 01 Subscribe to the MCP and enter your unique Scopus API credentials within Vinkius.
- 02 Connect this MCP in your preferred AI client (Claude, Cursor, etc.).
- 03 Ask your agent a question like, 'What were the citation trends for deep learning papers published last year?' and get instant results.

The bottom line is you get to treat the entire Scopus database like an extension of your natural conversation.

Built For

This MCP is for academics, bibliometricians, and data scientists who routinely analyze academic impact. If your job involves tracking research influence or auditing institutional output, this tool saves you days of manual database work.

Academic Researcher

Uses the MCP to quickly find relevant literature on a niche topic and track citation trends without manually exporting data sets.

Data Scientist (Knowledge Graph)

Integrates high-quality, structured metadata—like journal CiteScores or author affiliations—directly into knowledge graphs for analysis.

University Administrator

Audits institutional output and tracks the performance of faculty members by retrieving detailed author profiles and publication lists.

What Changes When You Connect

- 01 Stop manually exporting data sets. You can get a full abstract using the `get_abstract` tool, giving you all the metadata for one document immediately.

-
- 02** Build comprehensive author dossiers quickly. Use `search_author` to find profiles, then run `get_author` to pull detailed history and H-index counts.
-
- 03** Gauge research impact beyond simple citation counts. The MCP pulls altmetrics via `get_plumx_metrics`, showing how widely a paper is discussed online.
-
- 04** Compare journals accurately. Instead of reading vague rankings, use `get_serial_title` to pull concrete metrics like CiteScore and SJR for direct comparisons.
-
- 05** Track citation growth over time effortlessly. The `get_citation_overview` tool summarizes citations year by year, showing trends instead of just a single number.
-
- 06** Search with precision. Use the powerful `search_scopus` tool to query abstracts using Boolean syntax, drastically narrowing down your literature review results.
-

Real-World Applications

Assessing a rival's research focus

A researcher needs to know if their competitor is publishing in high-impact journals. They ask the agent to search for papers using `search_scopus` and then use `get_serial_title` on any resulting journal to check its CiteScore, giving them a clear competitive advantage.

Auditing an academic's portfolio

A university administrator wants to track a professor's career growth. They use `search_author` and then call `get_author` to retrieve the H-index, total publications count, and primary affiliation history instantly.

Writing a literature review chapter

A data scientist needs citation history for three key papers. They first run the `search_scopus` tool and then use `get_citation_overview` on each resulting DOI to generate year-by-year impact tables without leaving their agent environment.

Checking the reach of an early finding

A team has published a foundational paper. They ask the agent to run `get_plumx_metrics` on the document ID to see if it's getting buzz in media or social networks, which is crucial for grant applications.

Patterns to Avoid

Searching only by title

✗ AVOID

Trying to find all relevant papers on 'AI ethics' by just searching the title field. This misses context and related work.

✓ INSTEAD

Use ``search_scopus`` with Boolean operators, like `TITLE-ABS-KEY("AI ethics") AND PUBYEAR > 2020`, to narrow results across abstracts.

Assuming citation count is enough

✗ AVOID

Seeing a high total citation count but not knowing if that research was influential in the current field or decades ago.

✓ INSTEAD

Always use ``get_citation_overview`` to see citations broken down by year, revealing recent trends and impact.

Manually checking journal metrics

✗ AVOID

Opening three different journal websites just to compare their SJR or CiteScore metrics.

✓ INSTEAD

Use ``get_serial_title`` directly with the MCP. It pulls all key metrics like CiteScore and SJR into your agent's output in one step.

The Right Fit

You should use this MCP if your work requires systematic, quantitative analysis of academic knowledge—specifically bibliometrics or literature review support. If you need to find out *who* wrote something (author profiles), *where* they are affiliated (institutional data via `get_affiliation`), or the specific metrics of a publication (CiteScore from `get_serial_title`), this is your tool.

Don't use it if you just need general web information, current market trends, or primary source text that hasn't been indexed. If your goal is simply to find out 'what happened in the news today,' use a general search engine type tool instead. This MCP requires structured academic data, so keep your queries focused on authors, journals, and abstracts.

The Citation Treadmill: Collecting Research Metrics

Today, gathering research metrics is a tedious process. You jump between Scopus's portal, Google Scholar, and various journal websites. You copy the DOI from one site, paste it into another to check citation counts, then open a third tab to find out if the journal has an SJR score. This isn't analysis; it's glorified data entry.

With this MCP, you just tell your agent what you need. Ask for the citation overview for a specific DOI, or ask for the full metadata of an author's work. The MCP handles the API calls and metric aggregation in one go. You get structured, actionable data without opening another browser tab.

Scopus: Get Author Profiles and Impact Scores

The hardest part today is building a comprehensive view of an author's impact. You have to search by name, then find their institutional affiliation, then manually count the papers and track down the H-index for each one. It's slow, prone to error, and leaves you missing cross-referencing data.

Now, use this MCP. Start by running `search_author` on the person's name. The agent pulls all their academic details through `get_author`, giving you a clean, structured report that includes their current affiliation and H-index without any manual aggregation.

Scopus MCP: 10 Tools for Academic Research

These ten tools let you systematically examine document metadata, track author histories, compare journal quality, and calculate citation impact across the Scopus database.

#	TOOL	DESCRIPTION
01	<code>get_abstract</code>	Retrieves complete metadata for a single, specific academic document.
02	<code>get_affiliation</code>	Pulls detailed profile information about an educational or research institution.
03	<code>get_author</code>	Fetches a comprehensive professional profile for a named author, including their work history.
04	<code>get_citation_count</code>	Calculates the total number of times an abstract has been cited by other papers.
05	<code>get_citation_overview</code>	Generates a breakdown showing citation counts and summaries for specific years.
06	<code>get_plumx_metrics</code>	Retrieves altmetrics, measuring the document's visibility across social media and other platforms.
07	<code>get_serial_title</code>	Provides journal metrics like CiteScore, SJR, and SNIP to evaluate publication influence.
08	<code>search_affiliation</code>	Searches the database specifically for institutional profiles by name or keywords.
09	<code>search_author</code>	Locates and retrieves author profiles based on names, years, or affiliations.
10	<code>search_scopus</code>	Performs powerful searches across the entire database for abstracts using complex queries.

See It in Action

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

U Search Scopus for papers about 'Large Language Models' published in 2024.



I found several high-impact papers. The top results include 'A Survey of Large Language Models' and 'Efficient Training of LLMs'. Would you like me to fetch the abstract for any of these?

U Get the citation overview for DOI 10.1016/j.future.2023.01.001.



This document has been cited 45 times since its publication. In 2023 it received 12 citations, and so far in 2024, it has 33. It shows a strong upward trend in the field of cloud computing.

U Find the profile and H-index for author ID 7004212771.



Author ID 7004212771 corresponds to Dr. Jane Smith. She has an H-index of 52, with 142 total documents and over 8,500 citations. Her primary affiliation is Stanford University.

Frequently Asked Questions

01 How do I find citation trends using the Scopus MCP?

You use the `get_citation_overview` tool. Just provide a document's DOI or ID, and it generates a summary showing how many citations the paper received in specific years.

02 Can I search for papers published before 2010 with Scopus MCP?

Yes, you can. The `search_scopus` tool allows querying using historical fields like PUBYEAR and Boolean logic to target older datasets accurately.

03 What is the difference between getting abstract metadata and journal metrics via Scopus MCP?

Metadata (`get_abstract`) gives you info on one paper (title, authors). Metrics (`get_serial_title`) give you broad data about the publication source itself, like its overall CiteScore.

04 Do I need a specific key to use Scopus MCP?

Yes. You must subscribe and enter your valid Scopus API Key into Vinkius for the connection to work. The MCP requires this credential to access the data.

05 How do I find an institution's profile using Scopus MCP?







You can use `search_affiliation` first to locate the institution, then use `get_affiliation` to pull all detailed information about its history and structure.

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>"scopus": { "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

Scopus 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 Scopus. 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	Scopus MCP
Server ID	019e38e9-3754-73d9-baf3-a5aa0988e192
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/scopus.